

United States Court of Appeals for the Federal Circuit

COMCAST IP HOLDINGS I LLC,

Plaintiff – Cross-Appellant

v.

SPRINT COMMUNICATIONS COMPANY L.P.,

SPRINT SPECTRUM L.P.,

NEXTEL OPERATIONS, INC.

Defendants – Appellants

ON APPEAL FROM THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

Civil Action No. 1:12-205

Judge Richard G. Andrews

Brief of Defendants-Appellants

Rachelle H. Thompson
McGUIREWOODS LLP
434 Fayetteville Street
Raleigh, NC 27601
Tel: (919) 755-6572

Brian C. Riopelle
David E. Finkelson
Brian D. Schmalzbach
McGUIREWOODS LLP
800 East Canal Street
Richmond, VA 23219
Tel: (804) 775-1084

Counsel for Defendants-Appellants
Sprint Communications Company L.P.,
Sprint Spectrum L.P., and Nextel Operations, Inc.

April 4, 2016

UNITED STATES COURT OF APPEALS FOR THE FEDERAL CIRCUIT

Comcast IP Holdings I, LLC v. Sprint Communications Company

Case No. 15-1992

CERTIFICATE OF INTEREST

Counsel for the (petitioner) (appellant) (respondent) (appellee) (amicus) (name of party)

Appellant certifies the following (use "None" if applicable; use extra sheets if necessary):

1. The full name of every party or amicus represented by me is:

Sprint Communications Company L.P.

Sprint Spectrum L.P.

Nextel Operations, Inc.

2. The name of the real party in interest (Please only include any real party in interest NOT identified in Question 3. below) represented by me is:

None

3. All parent corporations and any publicly held companies that own 10 percent of the stock of the party or amicus curiae represented by me are listed below. (Please list each party or amicus curiae represented with the parent or publicly held company that owns 10 percent or more so they are distinguished separately.)

Sprint Communications Company L.P., Sprint Spectrum L.P., and Nextel Operations, Inc. are directly or indirectly owned by Sprint Communications, Inc. (SCI). SCI is a subsidiary of Sprint Corporation, a public company listed on the New York Stock Exchange. SoftBank Corp., a public company listed on the Tokyo Stock Exchange First Section, owns 10% or more of the stock of Sprint Corporation.

4. ☒ The names of all law firms and the partners or associates that appeared for the party or amicus now represented by me in the trial court or agency or are expected to appear in this court (and who have not or will not enter an appearance in this case) are:

McGuireWoods LLP: David Finkelson, Brian Riopelle, Brian Schmalzbach, Rachelle Thompson

Shook, Hardy & Bacon: Matthew Broaddus, Ryan Dykal, John Garretson, Aaron Hankel, Lynn Herndon,

Robert Reckers, Ryan Schletzbaum, Jared Tong, B. Trent Webb

Morris James LLP: Richard Herrmann, Mary Matterer

4/4/2016

Date

/s/ Brian D. Schmalzbach

Signature of counsel

Please Note: All questions must be answered

cc: _____

Brian D. Schmalzbach

Printed name of counsel

TABLE OF CONTENTS

	Page
INTRODUCTION	1
STATEMENT OF RELATED CASES	4
JURISDICTIONAL STATEMENT	4
STATEMENT OF THE ISSUES.....	4
STATEMENT OF THE CASE.....	6
I. The Low Patents.....	6
A. The '008 Patent	6
B. The '916 and '046 Patents	10
II. The Accused Call Flows	15
A. Call Flows Accused Of Infringing The '008 Patent	16
B. Call Flows Accused Of Infringing The '916 and '046 Patents	27
III. The District Court's Denial Of Judgment As A Matter Of Law.....	27
IV. The District Court's Award Of Prejudgment Interest.....	28
SUMMARY OF ARGUMENT	29
ARGUMENT	31
I. The District Court Invited The Jury To Ignore The Fundamental Nature Of The Asserted '008 Patent Claims By Misconstruing "Switched Telecommunications System.....	31
A. The '008 Patent Covers Using Resources On A Datagram-Based Network To Assist In Initiating A Call Through A Switched Telecommunications Network To A Call Destination—Not A Call Over A Datagram-Based Network.....	32
B. The District Court's Construction Fails To Account For The Requirement Of A Predetermined Bearer Channel	37
C. The District Court's Erroneous Construction Invited The Jury To Find Infringement Where Calls Were Transmitted To A Call Destination Over A Datagram- Based Network.....	39

TABLE OF CONTENTS

(continued)

	Page
II. There Is No Evidence To Support The Jury’s Evisceration Of The “Call Destination” Requirement And Related Terms in the ’008 Patent.....	41
A. Juries Must Apply The Ordinary Meaning Of Unconstrued Terms, Not Render Terms Superfluous.....	41
B. Defining “Call Destination” As “Any Point Along The Path Of A Call” Defies The Ordinary Meaning By Rendering The Terms Superfluous	43
C. The District Court Recognized That Comcast’s Theory Rendered “Call Destination” Superfluous But Tried To Resuscitate The Verdict Using Impermissible Post-Trial Claim Construction	46
III. The District Court Misconstrued The “Parsing” Requirement In The ’916 And ’046 Patents By Omitting The Complete Definition Of “Parsing” In The Prosecution History	50
A. The Prosecution History Reflects The Ordinary Meaning Of Claim Terms And Reflects Meanings Disclaimed By The Applicant.....	51
B. Any Proper Construction Of “Parsing” Must Distinguish “Predetermined Associations,” “Mapping,” “Lookup,” And “Translating” Because The File History Expressly Distinguishes Those Processes When Defining “Parsin.....	52
C. The District Court Failed To Give Effect To The Prosecution History.....	55
IV. The District Court Awarded Windfall Prejudgment Interest During A Time When Two Patents Could Not Have Been Infringed Because They Did Not Exist	57
A. Prejudgment Interest Is Available To Compensate During Periods Of Infringement	57
B. The District Court’s Award Overcompensates Comcast By Granting Interest Before There Could Have Been Infringement.....	59

TABLE OF CONTENTS
(continued)

	Page
C. There Is Ample Foundation In This Record To Calculate Prejudgment Interest Properly	63
CONCLUSION	64

ADDENDUM

Claim Construction Opinion [Docket 123].....	A1
Claim Construction Order [Docket 138]	A18
Order re: Claim Construction of the Term “Parsing” [Docket 298].....	A38
Judgment of October 22, 2014 [Docket 342]	A41
Judgment of November 14, 2014 [Docket 351]	A42
Opinion re: Post-Trial Motions [Docket 378].....	A48
Order re: Post-Trial Motions [Docket 379]	A70
U.S. Patent No. 7,012,916.....	A71
U.S. Patent No. 8,170,008.....	A108
U.S. Patent No. 8,204,046.....	A150

TABLE OF AUTHORITIES

	Page(s)
Cases	
<i>3M Innovative Props. Co. v. Tredegar Corp.</i> , 725 F.3d 1315 (Fed. Cir. 2013)	42
<i>Aguayo v. Universal Instruments Corp.</i> , 356 F. Supp. 2d 699 (S.D. Tex. 2005)	42
<i>Amgen Inc. v. F. Hoffman-La Roche, Ltd.</i> , 580 F.3d 1340 (Fed. Cir. 2009)	45
<i>Applied Med. Resources Corp. v. U.S. Surgical Corp.</i> , 435 F.3d 1356 (Fed. Cir. 2006)	59
<i>Biogen Idec, Inc. v. GlaxoSmithKline LLC</i> , 713 F.3d 1090 (Fed. Cir. 2013)	51, 53, 56
<i>CardSoft, LLC v. VeriFone, Inc.</i> , 807 F.3d 1346 (Fed. Cir. 2015)	31, 50
<i>DDR Holdings, LLC v. Hotels.com, L.P.</i> , 773 F.3d 1245 (Fed. Cir. 2014)	62, 63
<i>Ecolab, Inc. v. FMC Corp.</i> , 569 F.3d 1335 (Fed. Cir. 2009)	54
<i>Elekta Instrument S.A. v. O.U.R. Scientific Int’l, Inc.</i> , 214 F.3d 1302 (Fed. Cir. 2000)	42
<i>Epistar Corp. v. Int’l Trade Comm’n</i> , 566 F.3d 1321 (Fed. Cir. 2009)	56
<i>Fromson v. W. Litho Plate & Supply Co.</i> , 853 F.2d 1568 (Fed. Cir. 1988)	60, 61, 62
<i>Gen. Motors Corp. v. Devex Corp.</i> , 461 U.S. 648 (1983)	58
<i>Gustafson, Inc. v. Intersystems Indus. Prods., Inc.</i> , 897 F.2d 508 (Fed. Cir. 1990)	58, 62

<i>Haemonetics Corp. v. Baxter Healthcare Corp.</i> , 607 F.3d 776 (Fed. Cir. 2010)	39
<i>Hewlett-Packard Co. v. Mustek Sys., Inc.</i> , 340 F.3d 1314 (Fed. Cir. 2003)	47, 48
<i>LaserDynamics, Inc. v. Quanta Computer, Inc.</i> , 694 F.3d 51 (Fed. Cir. 2012)	58
<i>Lexion Med., LLC v. Northgate Tech., Inc.</i> , 641 F.3d 1352 (Fed. Cir. 2011)	32
<i>Lummus Indus., Inc. v. D.M. & E. Corp.</i> , 862 F.2d 267 (Fed. Cir. 1988)	63
<i>Marion v. TDI Inc.</i> , 591 F.3d 137 (3d Cir. 2010)	41
<i>N. Am. Container, Inc. v. Plastipak Packaging, Inc.</i> , 415 F.3d 1335 (Fed. Cir. 2005)	52
<i>Nickson Indus., Inc. v. Rol Mfg. Co., Ltd.</i> , 847 F.2d 795 (Fed. Cir. 1988)	58
<i>Oiness v. Walgreen Co.</i> , 88 F.3d 1025 (Fed. Cir. 1996)	58
<i>Omega Eng'g, Inc. v. Raytek Corp.</i> , 334 F.3d 1314 (Fed. Cir. 2003)	51, 55, 57
<i>Openwave Sys., Inc. v. Apple Inc.</i> , 808 F.3d 509 (Fed. Cir. 2015)	37
<i>Ormco Corp. v. Align Tech., Inc.</i> , 498 F.3d 1307 (Fed. Cir. 2007)	52
<i>Pannu v. Iolab Corp.</i> , 155 F.3d 1344 (Fed. Cir. 1998)	41, 45
<i>Perkin-Elmer Corp. v. Computervision Corp.</i> , 732 F.2d 888 (Fed. Cir. 1984)	41

<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005) (en banc)	34
<i>RFID Tracker, Ltd. v. Wal-Mart Stores, Inc.</i> , 342 F. App'x 628 (Fed. Cir. 2009)	52
<i>SafeTCare Mfg., Inc. v. Tele-Made, Inc.</i> , 497 F.3d 1262 (Fed. Cir. 2007)	37
<i>Shire Dev., LLC v. Watson Pharm., Inc.</i> , 787 F.3d 1359 (Fed. Cir. 2015)	56
<i>Sinclair Refining Co. v. Jenkins Petroleum Process Co.</i> , 289 U.S. 689 (1933)	61
<i>Smith & Nephew, Inc. v. Arthrex, Inc.</i> , 453 F. App'x 977 (Fed. Cir. 2011)	42
<i>Southwall Techs., Inc. v. Cardinal IG Co.</i> , 54 F.3d 1570 (Fed. Cir. 1995)	53
<i>Summit Tech., Inc. v. Nidek Co.</i> , 363 F.3d 1219 (Fed. Cir. 2004)	41
<i>Trs. of Columbia Univ. v. Symantec Corp.</i> , 811 F.3d 1359 (Fed. Cir. 2016)	34
<i>Uniloc USA v. Microsoft Corp.</i> , 632 F.3d 1292 (Fed. Cir. 2011)	42
<i>Uship Intellectual Props., LLC v. United States</i> , 714 F.3d 1311 (Fed. Cir. 2013)	57
<i>Vizio, Inc. v. Int'l Trade Comm'n</i> , 605 F.3d 1330 (Fed. Cir. 2010)	51
<i>Wi-Lan, Inc. v. Apple, Inc.</i> , 811 F.3d 455 (Fed. Cir. 2016)	48, 49
Statutes	
28 U.S.C. § 1295(a)(1)	4
28 U.S.C. § 1331	4

28 U.S.C. § 1338(a)	4
35 U.S.C. § 101	4
35 U.S.C. § 284	57, 61

INTRODUCTION

The claims at the heart of this patent infringement case are meant to create synergy between datagram-based networks like the Internet and traditional switched telephone networks. Specifically, the asserted claims of U.S. Patent No. 8,170,008 (“the ’008 patent”) address methods of accessing communication data from the world wide web (such as service resource items associated with a telephone number) over datagram-based networks for use in setting up bearer channels to transmit calls over switched networks. Through two critical errors, however, the district court transformed the asserted claims from a method of initiating or setting up a call over a *switched* network to a method of routing a call over a *datagram-based* network. That transformation effectively rewrote the claims and granted the plaintiff Comcast IP Holdings I LLC a windfall new monopoly over countless networks far beyond the intended scope of the claims.

The district court first erred by construing the term “switched telecommunication system” in a way that blurred the distinction between switched and datagram-based networks. Switched networks, as the ’008 patent teaches, rely on predetermined or dedicated bearer channels to transmit calls. Datagram (or packet-based) networks, by contrast, transmit data by independently routing data packets without following a predetermined bearer channel. The asserted claims require a predetermined bearer channel to transmit the call all the way from the

caller's telephone to the call destination. But by construing "switched telecommunication system" to include networks that use both predetermined bearer channels and independent routing, the district court invited the jury to find infringement in calls that travel over predominately packet-based networks that lack a predetermined bearer channel for the entire call flow. Comcast seized that invitation by arguing that the accused Sprint call flows infringed the '008 patent even though Comcast identified call destinations that were on datagram-based networks and could not be reached via predetermined bearer channels.

The district court compounded that claim construction error by sustaining a jury verdict that rendered the concept of a "call destination" superfluous. The asserted claims require a "call destination" that is determined by domain name system signaling. But Comcast's expert testimony and closing argument insisted that the call destination is not the final point on a call flow, but rather *any* point along the route, including any point along the datagram portion of the call flow. In post-trial proceedings, the court recognized after the fact that any theory with infinite call destinations has no meaningful call destination and thus reads the term out of the claims. Rather than granting judgment as a matter of law, the court substituted a new theory of its own making—that some points along the route of a call were call destinations, but not others. There was no evidence or any principled

basis for that theory, and the district court's attempt to rescue Comcast's verdict fails.

In addition to rewriting the '008 patent, the district court gave the jury an incomplete construction for two sister patents. U.S. Patent Nos. 7,012,916 ("the '916 patent") and 8,204,046 ("the '046 patent") recite the concept of "parsing" a number string. During the *ex parte* reexamination process, Comcast specified what "parsing" *is* ("a process of analyzing a string according to a set of rules of a grammar") and what it is *not*. In particular, Comcast distinguished parsing from several other processes, such as "mapping" or "translating." Comcast thus made clear that the ordinary meaning of "parsing" does not include those processes. The district court erred by not incorporating those negative limitations into its construction of "parsing."

The district court awarded Comcast one final windfall in the form of substantial prejudgment interest, which is meant to compensate for unpaid royalties during periods of infringement. The court's award, however, included interest for two patents over a six-year period when there could not have been infringement because the patents did not yet exist. Remand is necessary to correct all these errors.

STATEMENT OF RELATED CASES

No appeal from this civil action was previously before this or any other appellate court. Counsel for Sprint Communications Company L.P., Sprint Spectrum L.P., and Nextel Operations, Inc. (collectively “Sprint”) are not aware of any case pending in this or any other court that will directly affect or be directly affected by this Court’s decision.

JURISDICTIONAL STATEMENT

Comcast alleged a cause of action under the United States patent laws, 35 U.S.C. § 101 *et seq.*, and the district court exercised subject-matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a). On November 14, 2014, the district court entered a final judgment that Sprint infringes Claim 45 of the ’916 patent, Claims 90 and 113 of the ’046 patent, and Claims 1, 13, and 27 of the ’008 patent. A42-43. On November 19, 2014, Sprint filed a renewed motion for judgment as a matter of law under Rule 50(b) or, in the alternative, for a new trial under Rule 59. A251. The district court denied Sprint’s post-trial motion on August 10, 2015. A70. Sprint timely filed a notice of appeal on September 8, 2015. A6817. This Court has jurisdiction under 28 U.S.C. § 1295(a)(1).

STATEMENT OF THE ISSUES

- I. The asserted claims and specification of the ’008 patent require a predetermined bearer channel to transmit a call from the caller’s

- telephone to the call destination. Did the district court err by construing “switched telecommunication system” in a way that invited the jury to find infringement when a call travels on a packet-based network outside of a predetermined bearer channel and when the purported call destination is on that packet-based network?
- II. The only infringement theory presented to the jury for the ’008 patent required the jury to interpret “call destination” and related terms as “any point along the path of a call.” In sustaining the jury verdict, the district court posited a new construction of “call destination” as “some points along the path of a call, but not others.” Did the district court err by sustaining a jury verdict that rendered a critical claim term superfluous on the basis of a novel theory that was unsupported by evidence and never presented to the jury?
- III. During reexamination of the ’916 and ’046 patents, Comcast insisted on a definition of “parsing” that distinguished and disclaimed the concepts of “predetermined association,” “mapping,” “lookup,” and “translation.” Did the district court err by construing “parsing” without those negative limitations?
- IV. Courts may award prejudgment interest to compensate patent holders for the lost use of royalties when a patent is infringed. The district court

awarded Comcast prejudgment interest for multiple patents when those patents did not exist and could not have been infringed. Did the district court award Comcast windfall prejudgment interest?

STATEMENT OF THE CASE

In 2012, Comcast filed suit against Sprint, alleging that several varieties of Sprint telephone calls infringed eight total patents. A256. In 2013, Comcast withdrew allegations of infringement as to four of the asserted patents, leaving four patents-in-suit: U.S. Patent No. 6,873,694 (“the ‘694 patent”) and the ’008, ’916, and ’046 patents. A1572. The district court entered an order that the asserted claim of the ‘694 patent was invalid on July 16, 2014, A37, leaving the ’008, ’916, and ’046 patents, which are collectively referred to as the “Low patents” after the named inventor. The Low patents address methods of combining certain features of the Internet with traditional telephone networks.

I. The Low Patents

A. The ’008 Patent

1. The Asserted Claims

The ’008 patent, titled “Method and Apparatus for Accessing Communication Data Relevant to a Target Entity Identified by a Number String,” was issued on May 1, 2012. Comcast alleged that Sprint infringed the following

method claims, each of which addresses a “call through a switched telecommunication system”:

1. A method, comprising:

receiving, over a switched telecommunication system, a request;

determining, responsive to the request, a call destination using domain name system signaling; and

initiating a call through the switched telecommunication system between a calling party and the call destination that was determined as a result of said domain name system signaling.

13. A method, comprising:

receiving, over a switched telecommunication system, an indication of a called party;

determining, responsive to the indication of the called party, a call destination associated with the called party using domain name system signaling; and

initiating a call through the switched telecommunication system between a calling party and the call destination that was determined as a result of said domain name system signaling.

27. A method, comprising:

receiving a request from a first party;

determining by a computer, responsive to the request, an identifier of a second party using domain name system signaling;

setting up a call through the switched telecommunication system between the first party and the second party that was determined as a result of said domain name system signaling.

A146 (32:47-55); A147 (33:31-40, 34:39-47).

The specification describes the '008 patent as “a method of accessing service resource items that are intended to be used in setting up bearer channels through a switched telecommunications system.” A131 (1:27-30). It also states that “the term ‘switched telecommunication system’ means a system comprising a bearer network with switches for setting up a bearer channel through the network.” A131 (1:31-33). The specification further states that in the context of a switched telecommunication system, a “call” means “a communication through a bearer channel set up across the bearer network.” A131 (1:42-45). The term “communication system” has “a broader meaning” and “is intended to include datagram-based communication systems where each data packet is independently routed through a bearer network without following a predetermined bearer channel.” A131 (1:50-55).

2. The District Court’s Order Construing “Switched Telecommunication System”

The parties requested a *Markman* hearing for the term “switched telecommunication system” and proposed the following constructions:

Comcast	Sprint
A system comprising a bearer network with switches for setting up a bearer channel through the network. A datagram-based communication system where each data packet is independently routed through a bearer network without following a predetermined bearer channel is not a “switched telecommunication system.”	A system comprising a bearer network with switches for setting up a bearer channel through the network that does not include datagram-based communication systems where each data packet is independently routed through a bearer network without following a predetermined bearer channel.

The parties agreed that a network comprised of only datagram-based elements would not qualify as a switched telecommunication system. They disagreed, however, over whether a switched telecommunication system includes a network that contains some switched elements but over which a recited call destination is reached by datagram-based transmission that lacks the requisite predetermined bearer channel. Sprint’s construction excluded such networks by specifying that a switched telecommunication system “does not include datagram-based communication systems where each data packet is independently routed through a bearer network without following a predetermined bearer channel.” Comcast’s construction, by contrast, excluded only pure datagram-based systems.

The district court adopted Comcast’s construction. It stated, “The fact that the ‘communication system’ is understood to be broader than the ‘switched telecommunication system’ does not mean a ‘switched telecommunication system’ is precluded from having some overlap with elements of a ‘datagram-based

system.”” A4. The court concluded that “so long as a ‘switched telecommunication system’ indeed has switches, it may also have aspects of a ‘datagram-based system.’” A4.

B. The ’916 and ’046 Patents

1. The Asserted Claims

The Low ’916 and ’046 Patents both recite the term “parsing” as part of methods of accessing data for use in telephone services. The ’916 patent, also titled “Method and Apparatus for Accessing Communication Data Relevant To A Target Entity Identified By A Number String,” issued on March 14, 2006. A88 (1:1-4). Comcast alleged that Sprint infringed the independent method claim 45 of the ’916 patent:

45. A method of accessing communications data for contacting a target entity, said method comprising:

forming, from a number string identifying the target entity, a domain name by a process including *parsing* at least a substantial portion of the number string into at least a part of said domain name;

supplying the domain name formed to a DNS-type database system and receiving back a resource record including an URI for locating communications data associated with the domain name; and

using the URI received back to access said communications data.

A105 (36:56-67) (emphasis added).

Comcast also alleged infringement of independent method claims 90 and 113 of the '046 patent (“Method and Apparatus For Accessing Service Resource Items That Are For Use In A Telecommunications System”), issued June 19, 2012:

90. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including *parsing* at least a portion of the number string into at least a part of said domain name; and

supplying, by at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

113. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including *parsing* at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity; and

sending a message to the target entity identified using the URI.

A194 (45:65-46:6); A195 (48:30-40) (emphases added).

2. *Ex Parte* Reexamination Proceedings

In May and June of 2013, the United States Patent and Trademark Office (“Patent Office”) granted *ex parte* reexamination requests for the Low patents.¹ The Patent Office initially rejected the reexamined claims of the ’046 and ’916 patents, on prior art and other grounds. Following Comcast’s statements distinguishing prior art, however, the Patent Office issued *Ex Parte* Reexamination Certificates for the ’046 and ’916 patents. Reexamination proceedings for those patents concluded in February 2014.

During reexamination, Comcast provided an express definition of “parsing” as “a process of analyzing a string according to a set of rules.” A2310-11, 2531. Comcast’s expert Dr. Haas further described this process of “analyzing a string according to a set of rules” as analyzing an input string “according to the rules of a grammar.” A2531-32.

Comcast also distinguished “parsing” from “any predetermined association between the number string and a domain name.” A2531. Comcast further argued in the context of a related patent that “parsing” does not include mapping a telephone number to a domain name, A2868 (PTO interview summary for ’752

¹ The relevant reexaminations include those of the ’916 and ’046 patents, as well as related U.S. Patent Nos. 8,223,752 (“the ’752 patent”); 7,903,641; and 7,206,304 (“the ’304 patent”). All of these patents are in the same Low patent family, and all were originally asserted by Comcast in the district court. *See* A258, 260, 265.

Patent reflecting Comcast's argument that prior art "also lacks parsing and merely maps a telephone number to a domain name like a lookup table"), and does not include a lookup from a phone number to a domain name, A2959 (PTO interview summary for '046 Patent reflecting Comcast's argument that prior art "also does not provide parsing, as it is at best a mere lookup from a phone number to a domain name").² Finally, Comcast stated that "parsing" is "not the same as" the process of "translating" or "converting information from one form to another." A2327.

3. The District Court's Order Construing "Parsing"

The parties did not request construction of the term "parsing" at the initial *Markman* hearing. After reexamination proceedings, however, the district court granted Sprint's request for supplemental claim construction. The parties proposed the following constructions:

² Comcast ratified the statements in the Examiner's Interview Summary when it cited no inconsistencies or conflicting statements in the interview summary. *See* A2965 (Comcast "acknowledg[ing] the interview summary issued by the PTO on October 7, 2013, also summarizing that interview" and stating that "nothing in the PTO interview summary appears to be inconsistent . . .").

Comcast	Sprint
An automated process of determining the syntactic structure of a language unit by decomposing it into more elementary subunits and establishing the relationships among the subunits.	A process of analyzing a string according to a set of rules of a grammar. Parsing does not include mapping a telephone number to a domain name, nor a lookup from a telephone number to a domain name, and it does not refer to any predetermined association between a number string and a domain name. Parsing is also not the same as translating.

A38. At the supplemental *Markman* hearing, Comcast agreed with Sprint’s construction as “[a] process of analyzing a string according to a set of rules of a grammar,” but disputed the negative limitations for predetermined associations, mapping, lookups, and translating, A3180.

The district court adopted Comcast’s position from the hearing and construed “parsing” as “[a]n automated process of analyzing a string according to a set of rules of a grammar.” A40. The court did not dispute “that the patentee described things that are not by themselves parsing during the prosecution of the patent in order to elucidate the meaning of ‘parsing.’” A40. But the court did not think “that use of those negative limitations rose to the level of a clear disavowal of any claim scope contained within the plain meaning of parsing.” A40. Finally, the

court “note[d] that the cited [prior art] references are not in the record, and Sprint therefore has not shown any disclaimer.” A40.

II. The Accused Call Flows

At trial, Comcast accused the six different Sprint call flows below involving the Sprint Mobile Integration (“SMI”) service,³ Google Voice service,⁴ or the Airave 2 device⁵:

Accused Call Flow	Asserted Patent
When an SMI subscriber on a CDMA mobile handset ⁶ makes a call to any party except for a Sprint subscriber that is not an SMI or Google Voice user	'008
When a Google Voice subscriber makes a call to any party except for a Sprint subscriber that is not an SMI user	'008
When a Sprint subscriber on Sprint’s CDMA network makes a call to a user of an Airave 2 device	'008
When an SMI subscriber on a CDMA mobile handset makes a call to another SMI subscriber or a Google Voice subscriber	'916, '046
When a Google Voice subscriber makes a call to an SMI subscriber	'916, '046
When a Sprint subscriber using an Airave 2 device makes a call to an SMI subscriber or a Google Voice subscriber	'916, '046

³ The SMI service extends features available on workplace desk phones to mobile handsets. A4277.

⁴ The Google Voice service extends features available on Google to mobile handsets. A4278.

⁵ Sprint’s Airave 2 is a device that enhances indoor cellular coverage. A4279, 4705-06.

⁶ A CDMA handset is one that operates on Sprint’s wireless CDMA network. A4698-99.

None of the accused call flows travels exclusively over a switched network or predetermined bearer channels. Instead, each starts on Sprint's CDMA network, travels through Sprint's IP multimedia subsystem ("IMS") network (a packet-based network),⁷ and eventually returns to the CDMA network or the CDMA air interface for routing to the called handset.⁸

A. Call Flows Accused Of Infringing The '008 Patent

Comcast's expert, Dr. Burger, testified that three Sprint call flows infringed the asserted claims of the '008 patent, including the limitations of "determining . . . a call destination using domain name system signaling" and "determining . . . an identifier of a second party using domain name system signaling." The parties did not submit "call destination" or "identifier of a second party" for claim construction,⁹ so Comcast's expert, Dr. Burger, testified that a "call destination" is "the place that Sprint routes the call to," but repeatedly refused to testify that it is "the *last* place that the Sprint network routes a call to." A4542 (emphasis added). At closing argument, Comcast confirmed its theory that "a call destination can be any destination along the way in the route of a call." A5165; *see also* A5171-72

⁷ The IMS network includes an IMS core, which is a collection of servers on a packet-based network used to route and process signaling traffic for the accused Sprint voice calls. A4284, 4426, 4700.

⁸ The accused Airave 2-based call flows also traverse the packet-based Internet. A4371-73.

⁹ The terms "call destination" and "identifier of a second party" are only found in the claims. A4541.

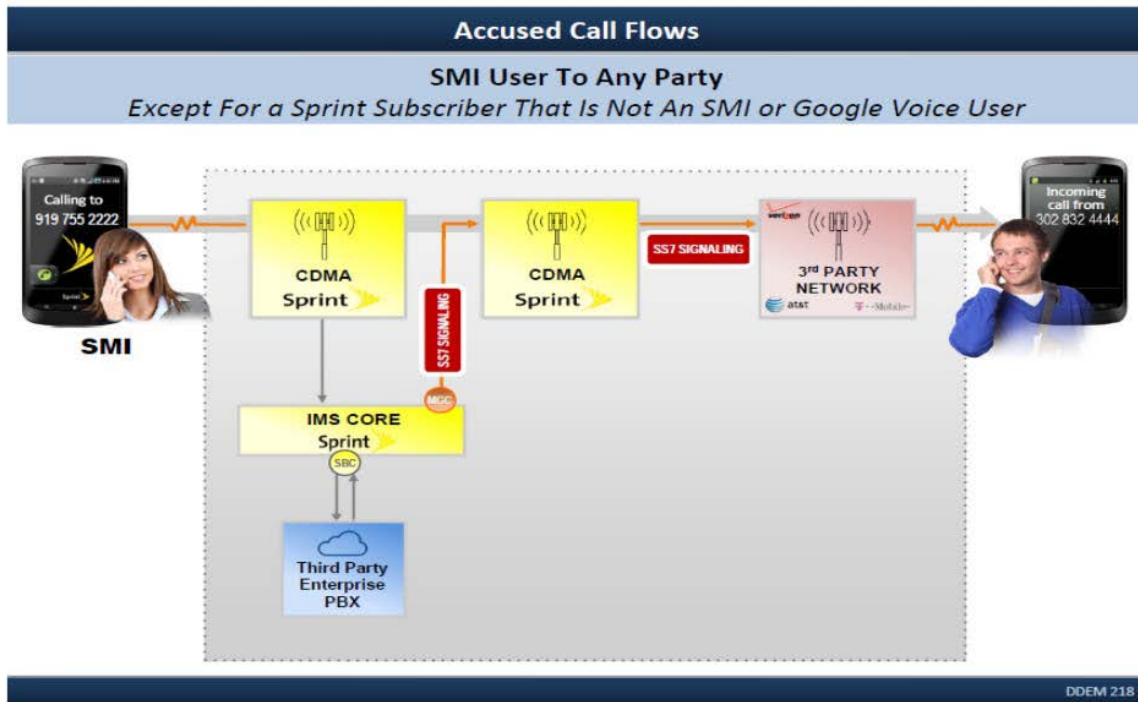
(“There are a multitude of destinations along the way when the call gets routed, and at some point there is domain name system signaling used.”). Dr. Burger offered the jury a similar definition of “identifier of a second party” as “how you get to the second party.” A4540.

1. When An SMI Subscriber On A CDMA Mobile Handset Makes A Call To Any Party Except For A Sprint Subscriber That Is Not An SMI Or Google Voice User

This accused call flow involves an SMI subscriber on a CDMA handset who makes a call to any party on a third party network, to an SMI subscriber, or to a Google Voice subscriber. Representative demonstratives are below.¹⁰ See A4391-96.

¹⁰ The accused call flow diagrams used for illustrative purposes in Section II.A were demonstrative exhibits and not entered into evidence at trial.

a. SMI Subscriber On A CDMA Handset Calls A Third Party Subscriber

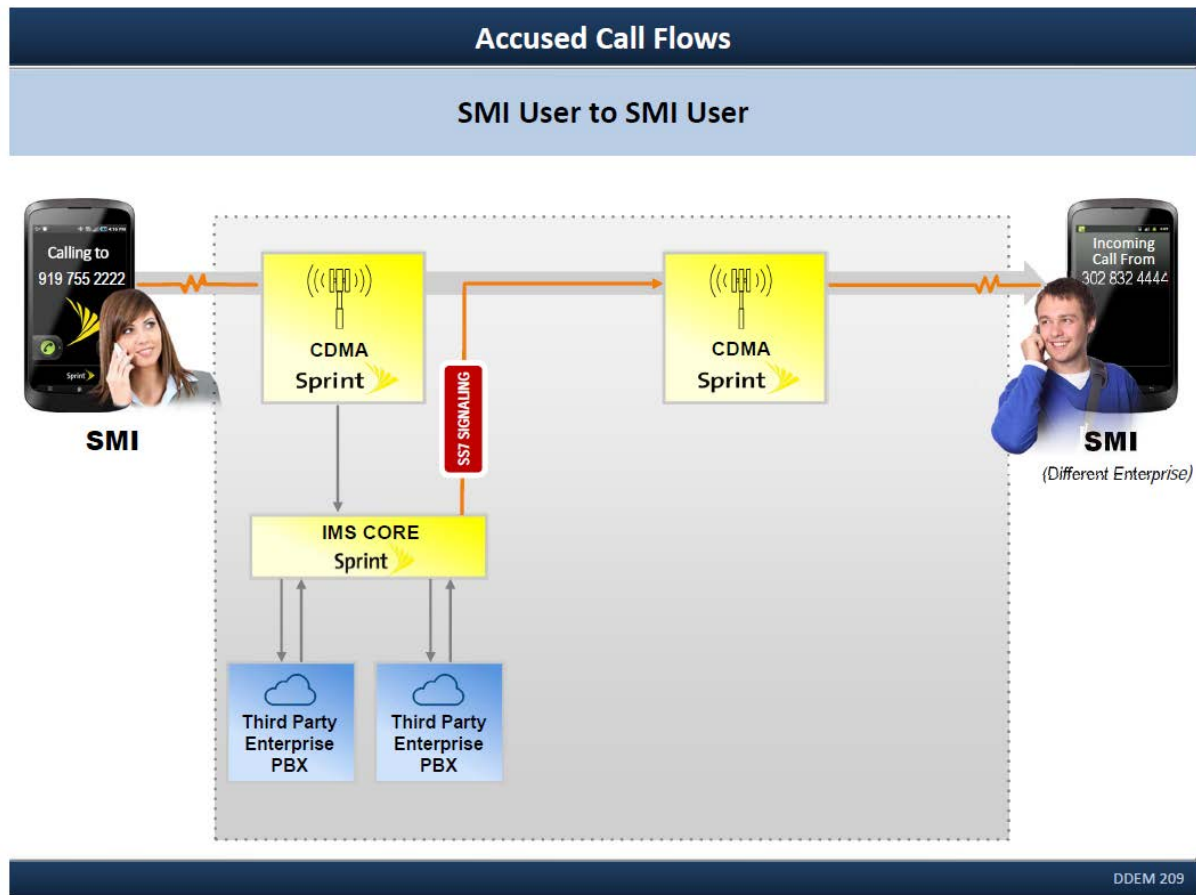


In this call flow, a call from an SMI subscriber on a mobile handset first travels on the switch-based CDMA network and then enters the packet-based IMS core via a media gateway controller, or “MGC.” A4342. The MGC is a server within Sprint’s IMS core that “handles the conversion of signalling from the CDMA network into the IP network” and “the conversion of the bearer channels” for the “switched voice on the CDMA side to the packet voice on the IMS side.” A4394-95, 4423. Once inside the IMS core, a network element named a call session control function (“CSCF”) uses domain name system signaling to determine an IP address used to route the call to a session border controller

(“SBC”). The SBC is another server on the IMS core. A4309-10, 4345-46, 4393. The SBC then routes the call to a Third Party Enterprise PBX outside of the Sprint network that is specific to the company to which the calling party belongs. A4345-46. That Third Party Enterprise PBX then sends the call back to the IMS core, where it is routed to another MGC that is responsible for a regional collection of area codes. A4346, 4348, 4395-96. The MGC is not specific to any particular telephone number or service provider. A4396. From the MGC, the call returns to the CDMA network and is ultimately routed to the called party’s handset on a third party network. A4342-43, 4350.

b. SMI Subscriber On A CDMA Handset Calls An SMI Subscriber

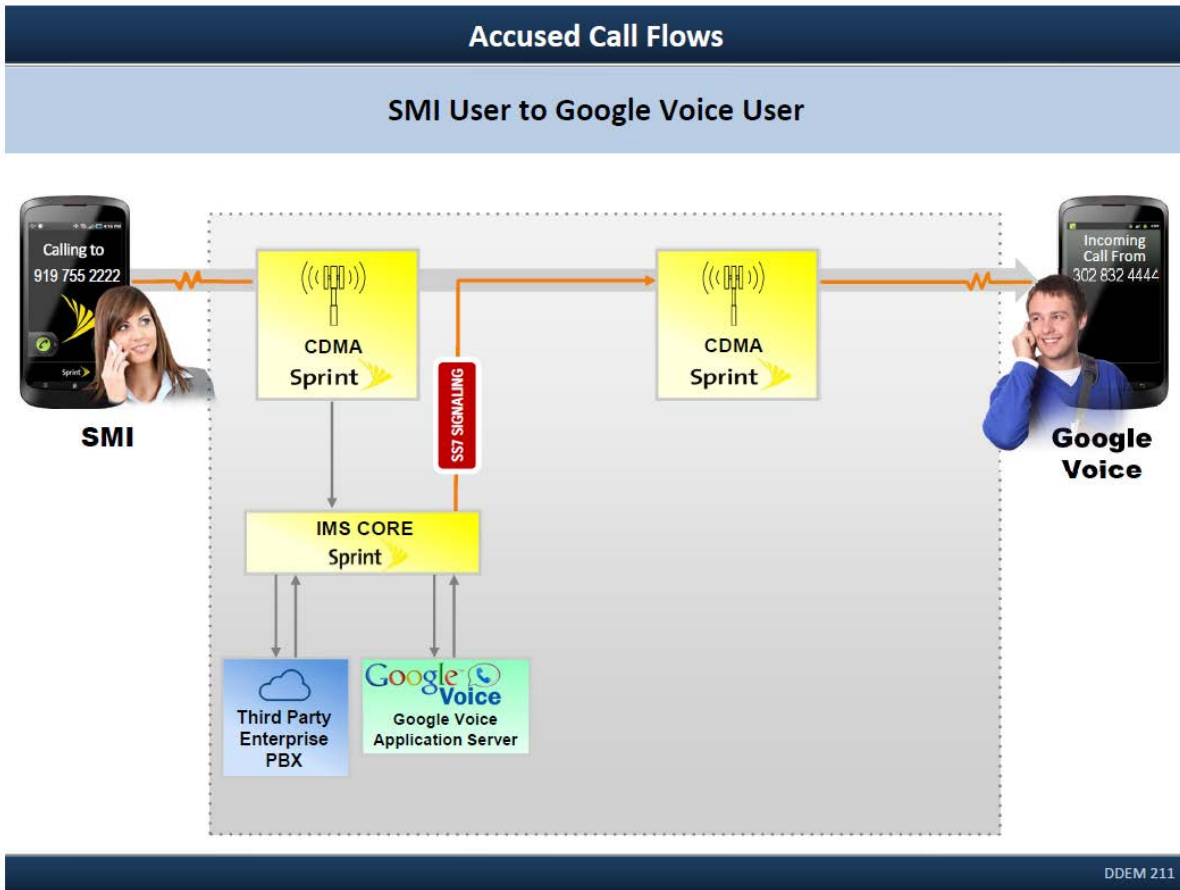
This accused call flow follows a path that is similar to that in Section A.1.a, except it travels from the IMS core to the Third Party Enterprise PBX twice.



Thus, the call travels from the CDMA network, to the IMS core, to a Third Party Enterprise PBX, back into the IMS core, to another Third Party Enterprise PBX, back into the IMS core, and back into the CDMA network for additional routing that terminates with a handset user that is an SMI subscriber. A4358-4365.

c. SMI Subscriber Calls A Google Voice Subscriber

This accused call flow follows a path that is similar to that in Section A.1.b, except one of the Third Party Enterprise PBXs has been replaced with a Google Voice Application Server.



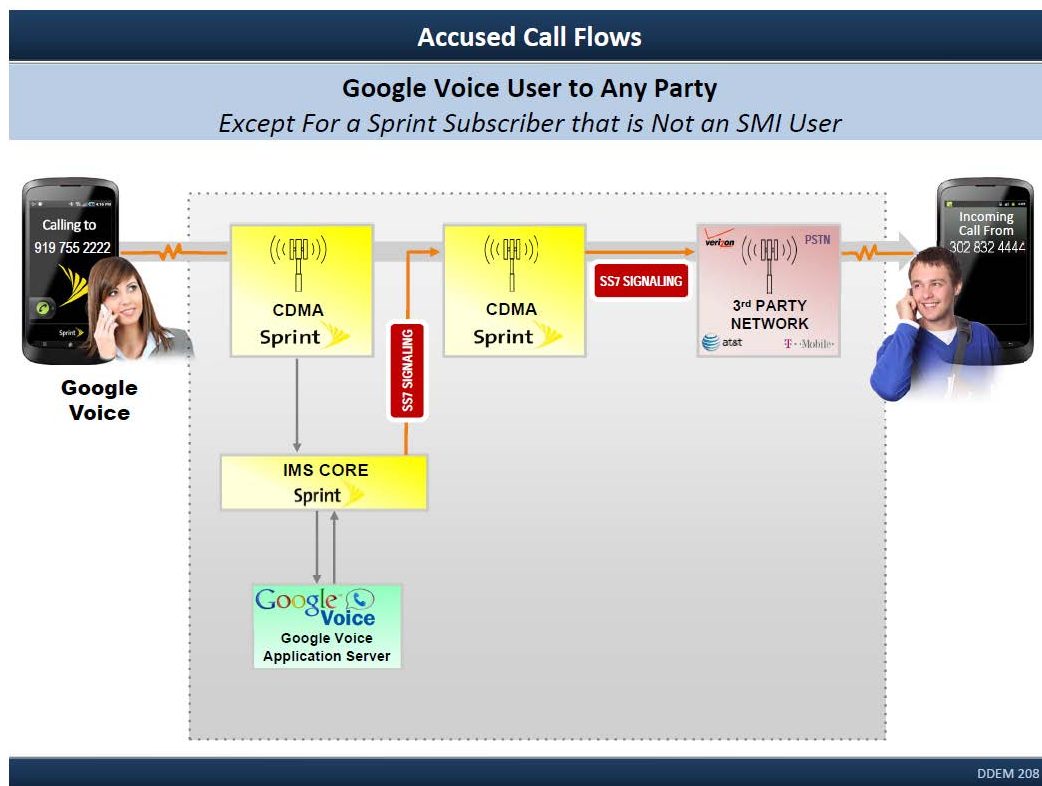
Thus, the call travels from the CDMA network, to the IMS core, to a Third Party Enterprise PBX, back into the IMS core, to a Google Voice Server, back into the IMS core, and back into the CDMA network for additional routing that terminates with a handset user that is a Google Voice subscriber. A4366-67.

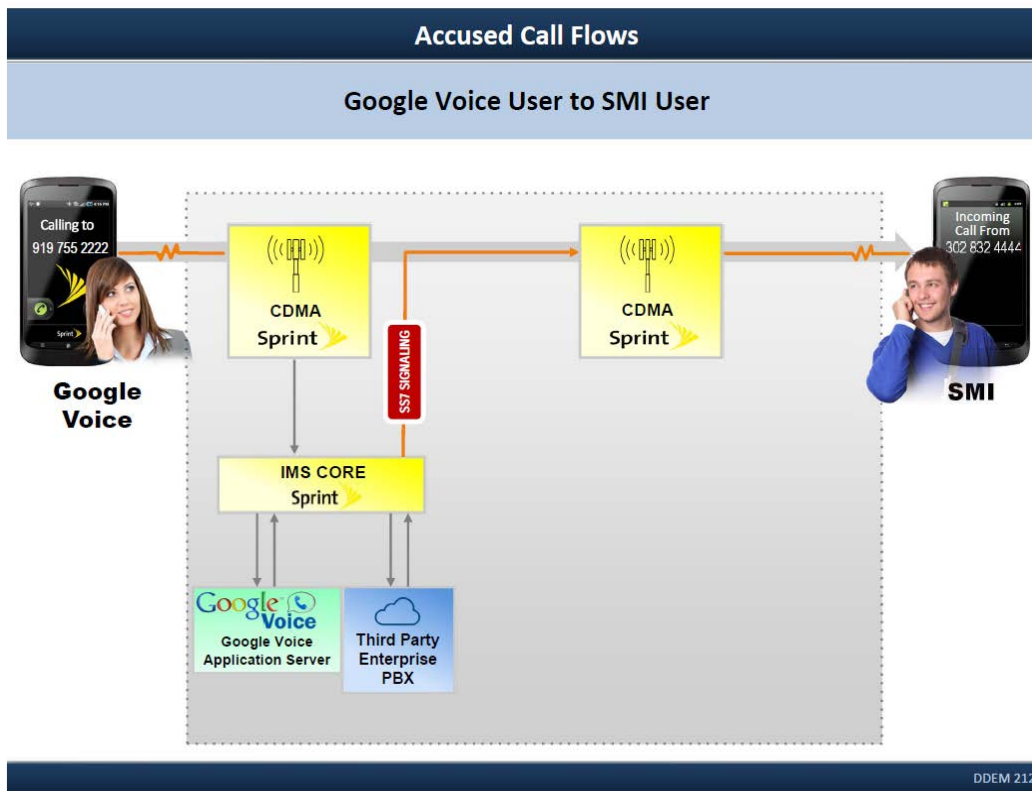
Dr. Burger opined that these call flows infringe the '008 patent. As to claim 1, he stated the call's first leg over the CDMA network, before it entered the packet-based IMS core, was a request received over a switched telecommunications system. A4439. He further stated that the MGC and SBC

servers on the IMS core were “call destinations” that satisfied the second limitation because that is where “the CSCF is determining: Where do I send this phone call to?” A4439. Finally, as to the third limitation requiring initiating a call through the switched telecommunication system between a calling party and a call destination, Dr. Burger opined that “the call is initiated through the switched telecommunication system because of the first part between a calling party and the call destination that was determined in that last step.” A4440. Dr. Burger concluded that the call flow infringed claims 13 and 27 for similar reasons. A4440-43; *see also* A4442 (opining that the “identifier of the second party” is an MGC or SBC).

2. When A Google Voice Subscriber Makes A Call To Any Party Except For A Sprint Subscriber That Is Not An SMI User

This accused call flow involves a Google Voice subscriber on a CDMA handset who makes a call to any party on a third party network or to an SMI subscriber, as shown in the demonstratives below.

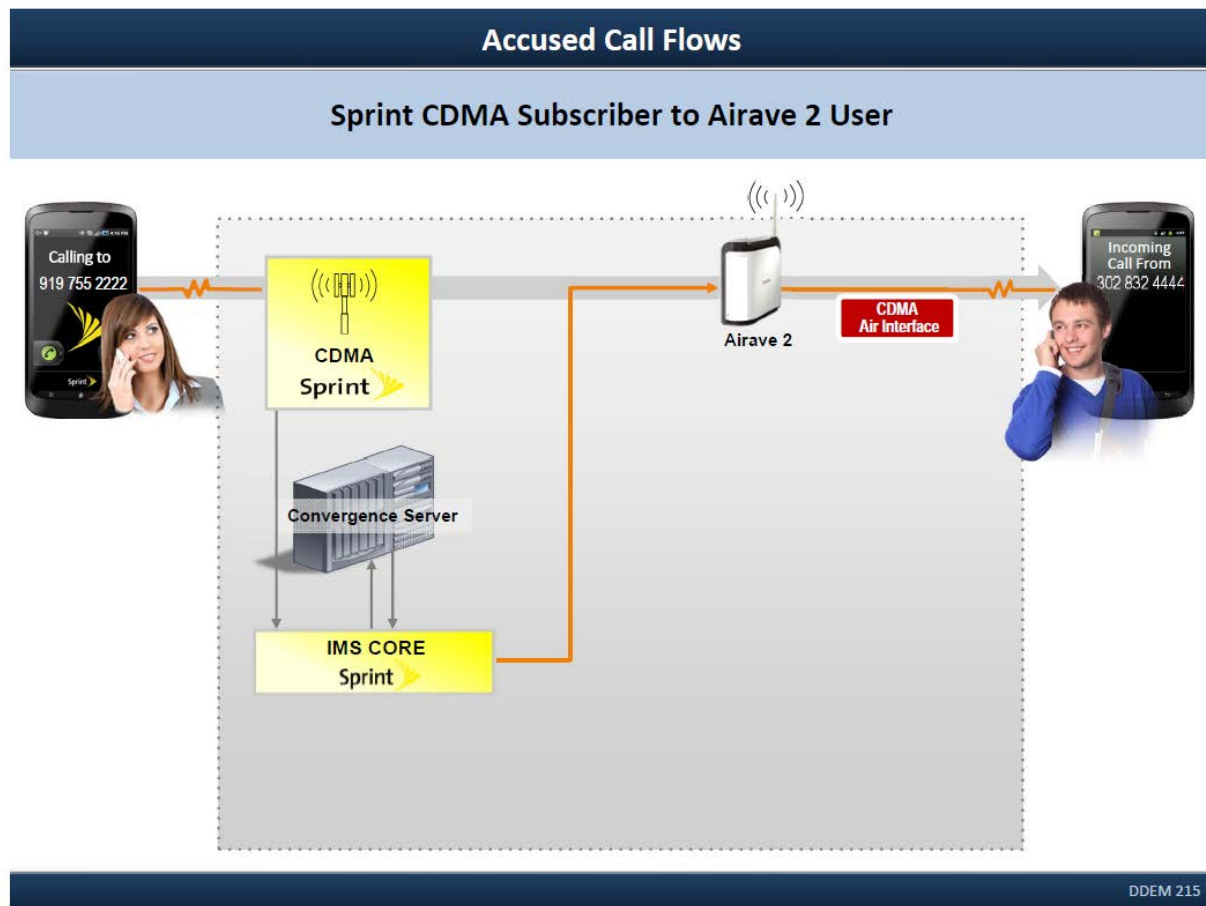




These call flows follow substantially the same steps as the previous call flows in Sections A.1.a-b. A4355. Instead of being routed from an SBC to a Third Party Enterprise PBX, the Google Voice calls are routed from an SBC to the Google Voice application server. A4355-56. Aside from that, those calls follow the same route as the calls from an SMI subscriber on a CDMA mobile handset who makes a call to any party except for a Sprint subscriber that is not an SMI or Google Voice user, and Dr. Burger offered similar infringement testimony. A4438-43.

3. When A Sprint Subscriber On Sprint's CDMA Network Makes A Call To A User Of An Airave 2 Device

As shown in the demonstrative below, this call travels from the CDMA network, to the IMS core, over the Internet to an Airave 2 device, and then over the CDMA air interface for routing to the called handset. A4371-73.



In this call flow, a call from a subscriber on Sprint's mobile network enters the packet-based IMS core. There, the CSCF uses domain name system signaling to determine the IP address of a convergence server and then routes the call to that server. A4375, 4377. The convergence server adds in the MAC address of the

Airave 2 device. A4375. Then, the CSCF routes the call to an SBC, which sends the call to the Airave 2 device over the Internet. A4373-76.¹¹ From there, the Airave 2 device determines which mobile handset to ring over the CDMA air interface. A4372-73, 4378.

Dr. Burger again opined that this call flow infringed the '008 patent. He stated that the first limitation in claim 1 was satisfied because the call traveled initially over Sprint's CDMA network. A4425-26. He concluded that the second limitation was satisfied because the IMS core determines the address of the convergence server by domain name system signaling, which in turn is used to identify the MAC address of the Airave 2 device. A4427. Dr. Burger concluded that the Airave 2 device was the recited call destination "because it's getting to that place which then can get to the subscriber." A4428. Finally, he stated that the third limitation was satisfied because "a call is initiated through the switch[ed] telecommunications system between the calling party . . . and the call destination that was determined, that MAC address of the Airave 2." A4429. Dr. Burger likewise concluded that the call flow infringed claims 13 and 27 of the '008 patent. A4431-37; *see also* A4442 (opining that the "identifier of the second party" is the MAC address of the Airave 2 device).

¹¹ Domain name system signaling is not used to locate the MAC address, SBC, or Airave 2 device. A4373, 4377.

B. Call Flows Accused Of Infringing The '916 And '046 Patents

Comcast accused three other call flows of infringing the '916 and '046 patents. A4443-44. As relevant here, Comcast alleged that each of those call flows practiced “parsing” because they used the ENUM method to route the calls through the IMS core. A4444, 4446. That method starts with a telephone number, removes any special characters, reverses the numbers, places a dot between each number, and adds “e.164.arpa” to the end. A4306, 4447. Dr. Burger opined that the ENUM method was “parsing” because “I’ve got this string of digits. And the analysis is read the digits, figure out the relationship between them. . . . Put them in a different order. And when you output them, put the dots in and the dots – 164.arpa.” A4453. Dr. Burger concluded that this constituted “an analysis pursuant to a set of rules.” A4453-54.

III. The District Court’s Denial Of Judgment As A Matter Of Law

Sprint moved for judgment as a matter of law at the close of Comcast’s case-in-chief and at the close of all the evidence. A4673-82, 5031-33, 5260. After the district court denied those motions, the jury found Sprint liable in the amount of \$7.5 million for infringing the '008, '916, and '046 patents. A41. The jury did not identify the amount of damages attributable to each asserted patent. A47. Sprint renewed its motion for judgment as a matter of law, as relevant here, because Comcast’s theory of “call destination” and “identifier of a second party” as “any

point along the path of a call” rendered those terms in the ’008 patent superfluous. A251.

The district court denied Sprint’s renewed motion for judgment as a matter of law. The court concluded that based on Dr. Burger’s testimony, a jury could find “several ‘call destinations.’” A54. The court stated, “To accept his testimony does not mean the term was read out of the claims; not every point along the route of a call is a call destination.” A54.

IV. The District Court’s Award Of Prejudgment Interest

After the jury awarded Comcast \$7.5 million in damages, Comcast sought prejudgment interest of \$1,691,640. A6710. Comcast sought interest based on a single hypothetical royalty negotiation in 2006, even though the ’046 and ’008 patents did not issue until 2012. Sprint argued that no prejudgment interest should be awarded for those patents during the six-year period when they did not exist and could not have been infringed. Sprint proposed to calculate prejudgment interest based on the relative values for each patent as calculated by Comcast’s damages expert, reduced in proportion to the amount that the jury reduced the expert’s total proposed damages. A6560.

The district court concluded that a “hypothetical negotiation in 2006 would include each of the asserted patents, even if they issued later.” A66. The court further concluded that it could not “reasonably infer from the jury’s verdict what

portion of the damages is attributed to each patent.” A66. The district court thus awarded the amount of prejudgment interest requested by Comcast. A67.

SUMMARY OF ARGUMENT

I. The district court construed the term “switched telecommunication system” in a way that failed to preserve the fundamental nature of the asserted claims. The ’008 patent draws a clear distinction between switched systems that rely on predetermined bearer channels to transmit a call to the call destination, and datagram-based systems in which the call data is packetized and independently routed to that call destination without following a predetermined bearer channel. The asserted claims and specification of the ’008 patent require a predetermined bearer channel between the calling party and call destination; that requirement cannot be satisfied if the call destination is in an independently routed datagram-based system. Instead of adopting Sprint’s construction that preserved the requirement of a predetermined bearer channel, the court adopted a construction that invited the jury to find a “switched telecommunication system” even when the accused system relies on primarily datagram-based transmission.

II. The district court sustained a jury verdict that rendered the concept of a “call destination” superfluous. The second limitation of two asserted ’008 patent claims requires a “call destination that was determined as a result of . . . domain name system signaling.” Through expert testimony and attorney argument,

however, Comcast invited the jury to interpret “call destination” as “any point along the route of a call.” The practical effect of that interpretation is that the second limitation will be satisfied whenever DNS signaling is used during a call, rendering the “call destination” requirement meaningless and sweeping countless networks under the patent. The district court recognized that Comcast’s theory would render “call destination” superfluous, but impermissibly substituted a novel and unsupported theory after trial that “call destination” means “some points along the route of a call, but not others.”

III. The district court misconstrued “parsing,” a contested term in the asserted claims of the ’916 and ’046 patents. During the *ex parte* reexamination of those patents, Comcast forcefully defined “parsing” in a successful attempt to stave off rejection. That definition included negative limitations for “predetermined association,” “mapping,” “lookup,” and “translating,” which Comcast disclaimed repeatedly. The district court erred in not distinguishing those relinquished concepts for the jury.

IV. The district court also granted Comcast a windfall by granting prejudgment interest for hypothetical royalties on two patents before those patents were issued. Prejudgment interest is meant to compensate patent owners who lost the use of unpaid royalties on infringed patents. But a patent cannot be infringed

before it is issued, and the district court abused its discretion by awarding interest for a six-year period when infringement was impossible.

ARGUMENT

I. The District Court Invited The Jury To Ignore The Fundamental Nature Of The Asserted '008 Patent Claims By Misconstruing “Switched Telecommunications System”

The asserted claims of the '008 patent are fundamentally directed to calls that travel on a predetermined bearer channel between the caller and call destination. The district court's first error was accepting a construction of the term “switched telecommunication system” that invited the jury to find infringement where a call traveled over a predominately *non*-switched system. In particular, the district court erred by adopting Comcast's confusing construction that was ripe for abuse rather than Sprint's clear and accurate construction that preserved the fundamental requirement of a bearer channel between caller and call destination. Because “the district court did not make any factual findings based on extrinsic evidence that underlie its construction[] of the disputed claim term,” this Court reviews its construction de novo. *CardSoft, LLC v. VeriFone, Inc.*, 807 F.3d 1346, 1350 (Fed. Cir. 2015).

The parties agree what a switched telecommunication system *is*: “a system comprising a bearer network with switches for setting up a bearer channel through the network.” Sprint's construction clarifies for the jury what a switched

telecommunication system is *not*: it “does not include datagram-based communication systems where each data packet is independently routed through a bearer network without following a predetermined bearer channel.” That is necessary to ensure what the claims require: a call that travels on a predetermined bearer channel spanning between the caller and call destination. Comcast’s construction, by contrast, embraces all datagram-based networks unless they include no switches whatsoever. The obvious—and intended—implication of Comcast’s construction is that a datagram-based system in which the accused call destination is reached by packet-based transmission is still deemed a “switched” system if it includes a single switch. But Comcast’s construction does not require the predetermined bearer channel that the claims require.

A. The ’008 Patent Covers Using Resources On A Datagram-Based Network To Assist In Initiating A Call Through A Switched Telecommunications Network To A Call Destination—Not A Call Over A Datagram-Based Network

The term “switched telecommunication system” must be construed to conform with its usage in the asserted claims. *See Lexion Med., LLC v. Northgate Tech., Inc.*, 641 F.3d 1352, 1356 (Fed. Cir. 2011) (“The customary meaning of a claim term is not determined in a vacuum and should be harmonized, to the extent possible, with the intrinsic record, as understood within the technological field of the invention.”). Here, that requires a construction that conforms with the

fundamental requirement of those claims: a call that travels over a predetermined bearer channel all the way from the caller to call destination.

Claims 1 and 13 illustrate the necessity of a predetermined bearer channel. As relevant here, those claims cover “initiating a call through the switched telecommunication system between a calling party and the call destination.” Based on the inventor’s own lexicography, *see* A131 (1:42-45), the parties and the district court agreed that “initiating a call through the switched telecommunication system” means “initiating a communication through a bearer channel set up across a bearer network of the switched telecommunication system.” A21. Thus claim 1 requires initiating a communication through a bearer channel set up across a bearer network between the calling party and the call destination. That “bearer channel . . . across a bearer network between the calling party and call destination” forms the claim’s basic requirement: a predetermined bearer channel on which the call travels from the calling party all the way to the call destination.¹²

¹² Claim 27 is substantially similar and likewise requires a call that is borne entirely on a predetermined bearer channel. Whereas claims 1 and 13 address “*initiating* a call through the switched telecommunication system between a *calling party and the call destination*, A146 (32:52-53), A147 (33:37-38), Claim 27 addresses “*setting up* a call through the switched telecommunication system between *the first party and the second party*.” A147 (34:44-45) (emphases added). Incorporating the agreed construction of “setting up a call,” Claim 27 requires “setting up a bearer channel through a bearer network of the switched telecommunication system between the first party and the second party.” Thus,

The '008 patent specification confirms that the claims require a predetermined bearer channel all the way from calling party to call destination. “[T]he specification is always highly relevant to the claim construction analysis. Usually, it is dispositive; it is the single best guide to the meaning of a disputed term.” *Phillips v. AWH Corp.*, 415 F.3d 1303, 1315 (Fed. Cir. 2005) (en banc).

First, the inventor’s lexicography confirms that requirement in the definition of “switched telecommunication system.” The inventor drew a crisp distinction between switched and datagram-based networks. The switched telecommunication system “compris[es] a bearer network with switches for setting up a bearer channel through the network.” A131 (1:32-33). But the term “communication system” has “a broader meaning” and “is intended to include datagram-based communication systems where each data packet is independently routed through a bearer network without following a predetermined bearer channel.” A131 (1:50-55). Thus, because the broader term “communication system” includes networks on which data need *not* follow a predetermined bearer channel to the destination, a switched telecommunication system is limited to networks on which the call *must* follow a predetermined bearer channel. *See Trs. of Columbia Univ. v. Symantec Corp.*, 811 F.3d 1359, 1364 (Fed. Cir. 2016) (“Even when guidance is not provided in explicit

Claim 27 also requires a predetermined bearer channel on which the call travels from the first party to the second party.

definitional format, the specification may define claim terms by implication such that the meaning may be found in or ascertained by a reading of the patent documents.”). To the extent a call on a network can pass outside a predetermined bearer channel and is independently routed to the destination, the network is necessarily a communications system and *not* a switched telecommunication system.

Second, the specification maintains the demarcation between switched- and datagram-based systems.¹³ The patent’s “Summary of the Invention” requires a “computer network” such as the Internet, A136 (11:66-12:1), that can locate “service resource items relating to setup control for bearer channels through [a switched] telecommunication system.” A135 (10:47-48). That computer network must be “logically distinct from the telecommunications system.” A136 (10:45-46). The purpose of that requirement is to exclude “computer networks that . . . effectively form part of the telecommunications system itself.” A136 (11:52-55). By excluding that scenario, the Summary of the Invention preserves the

¹³ Consistent with Sprint’s claim construction position, Comcast’s expert, Dr. Burger, considered these systems to be mutually exclusive when he testified that Sprint’s “CDMA network is a switched network, so it’s not a datagram-based one.” A4426.

requirement that a call travel on a predetermined bearer channel all the way from the calling party to the call destination.¹⁴

The specification reinforces the need for a call to travel on a predetermined bearer channel by distinguishing the invention from systems that allow calls to travel over both bearer channels and datagram-based networks. As one example of the prior art, which provided “no true interaction between the Internet and the [public switched telephone network, or PSTN],” the specification cites “dialup IP access to the Internet,” which involves the “setting up of a bearer channel over the PSTN for subsequently generated Internet traffic.” A135 (10:3-7). More importantly, the specification describes a form of Voice over Internet Protocol (VoIP), which entails “transferring digitized voice over the Internet to a service provider . . . [who] then connects into the local PSTN to access the desired phone and transfers across into the local PSTN the voice traffic received over the Internet.” A135 (10:12-16). Both of those are “a simple chaining together of the Internet and PSTN” in which calls flowed over a datagram network and a switched network consecutively with “no true interaction between the Internet and the

¹⁴ The Summary of the Invention notes that one type of computer network that is not “logically distinct” from the switched telecommunication system is a network “dedicated to the management or monitoring of the bearer network.” A136 (11:53-54). If a monitoring network is excluded because it “effectively form[s] part of the telecommunications system itself,” so much more is a computer network that actually bears the call. A136 (11:54-55).

PSTN.” A135 (10:6-7, 22); *see also Openwave Sys., Inc. v. Apple Inc.*, 808 F.3d 509, 513 (Fed. Cir. 2015) (“repeated derogatory statements about [a particular embodiment] reasonably may be viewed as a disavowal”). The invention, by contrast, is meant to “facilitate[] the integration of the PSTN and the [Internet]” by using a datagram-based network to access resources for “setting up bearer channels through a switched telecommunication system” rather than using a datagram-based network itself for bearing the call to the call destination. A135 (10:33-41); *SafeTCare Mfg., Inc. v. Tele-Made, Inc.*, 497 F.3d 1262, 1270 (Fed. Cir. 2007) (finding disclaimer where the specification repeatedly indicated that the invention operated by “pushing (as opposed to pulling) forces,” and then characterized the “pushing forces” as “an important feature of the present invention”).

B. The District Court’s Construction Fails To Account For The Requirement Of A Predetermined Bearer Channel

Because the claims, on their own and read in light of the specification, require a predetermined bearer channel from calling party to call destination,¹⁵ the district court should have construed “switched telecommunication system” consistent with that requirement. Its failure to do so requires reversal.

¹⁵ The requirement of a predetermined bearer channel described in Sections I.B-C applies with equal force to a call between a “first party” and a “second party” in Claim 27.

Comcast's construction contained a loophole that eviscerated the requirement of a predetermined bearer channel. In particular, Comcast's construction swept in all datagram-based networks so long as they are coupled to a switch. *See* A1104 (asserting that a switched telecommunications system includes "any system where *both* sets of elements can be found—i.e., (1) a bearer network with switches for setting up a bearer channel through the network, *and* (2) datagram packets that are independently routed through a bearer network without following a predetermined bearer channel."). The problem with the hybrid systems incorporated into Comcast's construction is that they do not provide a bearer channel through a bearer network between the calling party and the call destination. Yet by excluding only pure datagram systems, Comcast's construction invited the jury to deem even predominately datagram-based networks as switched.

The district court agreed with Comcast's construction on the grounds that a "system with elements of both switches and a 'datagram-based system' is not necessarily outside the scope of a 'switched telecommunication system.'" A4 (stating that a switched telecommunication system is not "precluded from having some overlap with elements of a 'datagram-based system'"). That observation is irrelevant. Of course there can be an overlap between switched and datagram systems: the '008 patent itself is meant to take advantage of a particular kind of overlap, one in which service resource items accessed over a datagram-based

network are used for setting up bearer channels over the switched network. A135 (10:39-41). But in no case do the asserted '008 patent claims allow that a datagram network actually bear the call between caller and call destination. The district court erred in defining “switched telecommunication system” to include a system that does not transmit a call to the call destination on a predetermined bearer channel.

C. The District Court’s Erroneous Construction Invited The Jury To Find Infringement Where Calls Were Transmitted To A Call Destination Over A Datagram-Based Network

Comcast’s construction was not only confusing in the abstract, but it successfully confused the jury into finding infringement where calls traveled over a datagram network and the call destination itself was actually on a datagram-based network. The verdict must be vacated. *See Haemonetics Corp. v. Baxter Healthcare Corp.*, 607 F.3d 776, 784 (Fed. Cir. 2010) (vacating a jury verdict of infringement that relied on the district court’s incorrect claim construction).

Comcast presented no evidence that any of the accused call flows traveled exclusively on the CDMA switched network used by Sprint. *See* A4426 (“The CDMA network is a switched network, so it’s not a datagram-based one, and, in fact, it is a system comprising a bearer network, that’s all the trunks that . . . set up the switch through that network.”). To the contrary, Comcast’s expert testified that each of the accused call flows traveled in Sprint’s IMS core, a datagram or packet-

based network. *See* A4426 (describing the IMS as a “packet-based network”); A4700 (Sprint’s expert describing Sprint’s IMS network as “a subcomponent of what our IP network or our packet data network is”). In fact, each of the purported call destinations identified by Comcast’s expert was in the IMS core or reached on the Internet (the quintessential datagram-based network). A4373, 4424.

For example, for the call flow from a Sprint CDMA handset to a user of an Airave 2 device, Comcast’s expert told the jury that the “switched telecommunication system” limitation was satisfied because the call is first sent from the handset to a mobile switching center on the CDMA network. A4425. But the expert identified the Airave 2 device as the call destination, A4427, and undisputed evidence showed that calls reach the device by traveling through the datagram-based IMS core and then through the Internet, A4373. Thus Comcast’s expert persuaded the jury to find a “switched telecommunication system” even though the bearer channel ceased to bear the call after the first step of the call flow.

Likewise, the expert told the jury that the “switched telecommunication system” limitation was satisfied for calls from SMI or Google Voice users on CDMA handsets to any party except for a Sprint subscriber that is not an SMI and/or Google Voice subscriber because, again, the call is first sent to a mobile switching center. A4439. But the expert identified an SBC or MGC in the IMS core as the call destination. Because the IMS core is a datagram-based network,

there can be no bearer channel stretching from the calling party to the call destination in the IMS core, as identified by Comcast.¹⁶

Comcast's trial evidence thus disproved the existence of any predetermined bearer channel from the caller to any accused call destination. Because Comcast's trial argument exploited the overbreadth baked into its faulty claim construction, the claim construction order must be reversed and the jury verdict vacated.

II. There Is No Evidence To Support The Jury's Evisceration Of The "Call Destination" Requirement And Related Terms in the '008 Patent

A. Juries Must Apply The Ordinary Meaning Of Unconstrued Terms, Not Render Terms Superfluous

The denial of a motion for judgment as a matter of law (JMOL) is reviewed *de novo*. *See Summit Tech., Inc. v. Nidek Co.*, 363 F.3d 1219, 1223 (Fed. Cir. 2004); *Marion v. TDI Inc.*, 591 F.3d 137, 146 (3d Cir. 2010). JMOL is proper when "the jury's findings, presumed or express, are not supported by substantial evidence or, if they were, that the legal conclusion implied [by] the jury's verdict cannot in law be supported by those findings." *Pannu v. Iolab Corp.*, 155 F.3d 1344, 1348 (Fed. Cir. 1998); *see also Perkin-Elmer Corp. v. Computervision Corp.*, 732 F.2d 888, 893 (Fed. Cir. 1984). As the plaintiff, Comcast bears the

¹⁶ Comcast's identification of "call destination[s]" is at odds with the concept of a "call," which, in the context of a switched telecommunication system, refers to "a communication *through* a bearer channel set up across the bearer network." A131 (1:42-45) (emphasis added).

burden of proving that Sprint practiced each element of the asserted claims. *Uniloc USA v. Microsoft Corp.*, 632 F.3d 1292, 1301 (Fed. Cir. 2011).

Because no party asked it to construe the terms “call destination” or “identifier of a second party,” the district court instructed the jury to apply the “common meaning” of those terms. A5105. “Common meaning” means that the terms must have *meaning*; an interpretation that renders a term superfluous does not reflect its common meaning. *See 3M Innovative Props. Co. v. Tredegar Corp.*, 725 F.3d 1315, 1329 (Fed. Cir. 2013) (claim terms “should be understood in a way that does not render the actual words of the claim superfluous”); *Elekta Instrument S.A. v. O.U.R. Scientific Int’l, Inc.*, 214 F.3d 1302, 1306-07 (Fed. Cir. 2000) (holding that the limitation “only within a zone extending between latitudes 30°-45°” is not satisfied on a device with radiation sources extending between 14° and 43° because “[a]ny other conclusion renders the reference to 30° superfluous”); *see also* Antonin Scalia & Bryan A. Garner, *Reading Law* 174 (describing the canon against superfluity as “true not just of legal texts but of all sensible writing”). This Court has scrutinized jury verdicts to ensure that juries do not depart from the common meaning of claim terms. *See Smith & Nephew, Inc. v. Arthrex, Inc.*, 453 F. App’x 977, 980 (Fed. Cir. 2011) (reversing the denial of JMOL where the plaintiff’s “theory of infringement is inconsistent with the plain language” of the asserted claims); *see also Aguayo v. Universal Instruments Corp.*, 356 F. Supp. 2d

699, 729-30 (S.D. Tex. 2005) (granting JMOL of noninfringement where the plaintiff's theory of the unconstrued term "indication" required additional steps beyond those required by the plain meaning).

B. Defining "Call Destination" And Related Terms As "Any Point Along The Path Of A Call" Defies The Ordinary Meaning By Rendering The Terms Superfluous

Claims 1 and 13 of the '008 patent require a "call destination that was determined as a result of . . . domain name system signaling." Although Sprint presented evidence that "call destination" referred to the final point on the path of a call, *i.e.* the called party, Comcast persuaded the jury that "call destination" means *any* point on the path of a call. The problem with Comcast's theory is that if the call destination is anywhere along a call flow, then the "call destination" limitation is meaningless: as long as a call is initiated and DNS signaling is used at some point during the call flow, the call destination requirement will always be satisfied. Comcast's theory thus renders a critical element of the asserted claims superfluous, deprives "call destination" of *any* meaning, and threatens to sweep countless networks that happen to use DNS signaling under the '008 patent.

Comcast's expert, Dr. Burger, testified that the call destination is wherever Sprint routes its calls to:

Q. And your infringement opinions on claims 1 and 13 of the '008 patent are based on your opinion that the call destination is the last place that the Sprint network routes a call to; is that correct?

A. Actually, it's the place that Sprint routes the calls to.

A4542. Dr. Burger provided no limiting principle for what could be considered the call destination, as long as the so-called destination is part of the call flow. Thus, he told the jury that mere waypoints along the call route, such as the MGC and SBC servers on the IMS core, were "call destinations" because that is where "the CSCF is determining: Where do I send this phone call to?" A4439; *see also* A4428 ("[The Airave 2 device] is [the] call destination because it's getting to that place which then can get to the subscriber.").

Comcast's closing argument cemented Dr. Burger's testimony that the call destination was anywhere along the route of the call: "We say that a call destination can be any destination along the way in the route of a call." A5165.

Comcast elaborated:

So the position Comcast takes in this case is that when it says there is a call destination using domain name system signaling, that means it can be a multitude of destinations. And in this case, that is exactly what we're talking about. There are a multitude of destinations along the way when the call gets routed, and at some point there is domain name system signaling used.

A5171-72. Thus Comcast explained to the jury that the call destination is any routing point along a call flow, and offered no evidence that any point along that call flow would not be considered a call destination.

Comcast's trial theory likewise trivialized the term "identifier of a second party" in claim 27 of the '008 patent. At trial, Dr. Burger extended Comcast's "any point along the way" theory to "identifier of a second party" by testifying that the term means only "how you get to the second party." A4540. Under Dr. Burger's theory, any piece of equipment, or its associated address, anywhere in a given call flow would qualify as "an identifier of a second party," even if it does not identify the second party at all. For example, Comcast identified an MGC as an "identifier of a second party" (i.e., a service provider), A4442, 4537, despite un rebutted evidence that MGCs do not identify and are not specific to any dialed party or even to any service provider. A4396. Thus, Comcast's interpretation renders "identifier of a second party" superfluous because as long as DNS signaling is used, there will be an identifier of a second party somewhere along the call flow.

Because Comcast's only theory of infringement rendered the terms "call destination" and "identifier of a second party" superfluous, the legal premise of the jury's verdict is incorrect. *See Pannu*, 155 F.3d at 1348 (JMOL is proper where the "legal conclusion implied [by] the jury's verdict cannot in law be supported by those findings"). And because Comcast's only theory for satisfying each of those terms reads the terms out of the claim, Comcast failed as a matter of law to meet its burden of proving that every element of the asserted claims is met. *See Amgen Inc.*

v. F. Hoffman-La Roche, Ltd., 580 F.3d 1340, 1374 (Fed. Cir. 2009) (there is no literal infringement as a matter of law if any claim limitation is absent from the accused device).

C. The District Court Recognized That Comcast’s Theory Rendered “Call Destination” Superfluous But Tried To Resuscitate The Verdict Using Impermissible Post-Trial Claim Construction

For the reasons above, the jury could not have applied any common meaning of “call destination” because Comcast’s “any point along the way” theory read that term out of the claims altogether. In fact, the district court recognized that if every point along the route of a call was a “call destination,” then that term would be rendered superfluous.¹⁷ But rather than granting JMOL, the court, in an attempt to cure the jury’s unsupportable verdict, crafted a new theory of its own that was never presented to the jury. Instead of Comcast’s “any point along the way” theory, the district court substituted a “*some* points along the way” theory: “not every point along the route of a call is a call destination.” A54. Substituting that theory was error for two reasons: the court never gave the jury a chance to consider the novel construction it applied to the claims, and the construction contradicted the evidence actually presented to the jury.

¹⁷ The term “identifier of a second party” would likewise be rendered superfluous under the district court’s reasoning.

First, the district court erred by resorting to a novel post-trial claim construction in an attempt to rescue the flawed jury verdict—an approach this Court soundly rejected in *Hewlett-Packard Co. v. Mustek Sys., Inc.*, 340 F.3d 1314, 1321 (Fed. Cir. 2003). There, the independent claim was drawn to “[a]n optical scanner device for producing machine-readable data representative of a color image of a scanned object comprising . . . scan speed indicating means for generating a scan speed signal indicating a selected one of different scan speeds.” *Id.* at 1319 n.2. The district court construed “scan speed indicating means for generating a scan speed signal” as “a scan speed selector which, *based on the user’s selection*, generates a scan speed signal that indicates a selected one of different scan speeds” or the equivalent. *Id.* at 1319. Although the district court gave the jury no further instructions regarding “based on the user’s selection,” after the jury found infringement, the court further construed the claim to lack a “‘one-to-one correspondence requirement’ between the variable adjusted by the scanner user and the scan speed” or any “requirement that speed be measured or indicated in any particular units.” *Id.* at 1320. “The claims so construed,” the district court concluded that the accused device’s ability to select *resolution* was sufficiently related to the ability to select scan speed to justify the verdict. *Id.*

This Court held that the district court could not rescue the jury verdict with that post-hoc construction. “[W]here the parties and the district court elect to

provide the jury only with the claim language itself, . . . it is too late at the JMOL stage to argue for or adopt a new and more detailed interpretation of the claim language and test the jury verdict by that new and more detailed interpretation.” *Id.* at 1321. Instead, “[t]he verdict must be tested by the charge actually given and by giving the ordinary meaning of the language of the jury instruction.” *Id.* Because the charge actually given did not permit the jury to find infringement based on a *resolution* selector, this Court vacated the verdict.

Here, the district court likewise gave an impermissible post-trial construction of “call destination.” Although the court instructed the jury to apply the common meaning of that term, the court itself went beyond “the charge actually given,” *Hewlett-Packard*, 340 F.3d at 1321, by concluding that there can be multiple call destinations but that “not every point along the route of a call is a call destination.” A54. It is clear that the court tested the jury’s verdict against that new construction because if it had not, and every point along the route of a call *were* a call destination, then by clear implication the term would be “read out of the claims.” *Id.* However, it is “too late at the JMOL stage to . . . adopt a new and more detailed interpretation of the claim language and test the jury verdict by that new and more detailed interpretation.” *Hewlett Packard*, 340 F.3d at 1321; *Wi-Lan, Inc. v. Apple, Inc.*, 811 F.3d 455, 465 (Fed. Cir. 2016) (“At the JMOL stage,

the question for the trial court is limited to whether substantial evidence supports the jury's verdict under the issued construction."").

The district court's novel post-trial construction was also improper because it had no foundation in the trial record. The court cited no evidentiary support. And, for the reasons above, Comcast's trial evidence does not recognize any boundaries on what could be considered the call destination within the call flow because Comcast proffered that a call destination is wherever Sprint routes the calls, even merely from one point in the network to another. A4542. Furthermore, Comcast's closing argument dispelled any limits on what constitutes the call destination: "We say that a call destination can be any destination along the way in the route of a call." A5165. Thus the only evidence and argument in support of the verdict reached by the jury contradict the district court's "not every point along the way" theory. *See Wi-Lan*, 811 F.3d at 465-66 (rejecting a post-trial construction as inconsistent with the expert testimony).

Even if there were any reason to consider some points along a call flow to be call destinations—but not others—the district court never explained how to determine which points are call destinations and which are not. And without a principle for divining call destinations from mere points along the way, there is no basis for concluding that the purported call destinations identified by Comcast are in fact call destinations. There is accordingly no basis for sustaining a jury verdict

necessarily founded on evidence and argument that contradict the district court's novel theory.

III. The District Court Misconstrued The "Parsing" Requirement In The '916 And '046 Patents By Omitting The Complete Definition Of "Parsing" In The Prosecution History

The district court made another claim construction error when it gave the jury an incomplete definition of "parsing," a critical term in the '916 and '046 patents. This Court reviews that error de novo. *CardSoft, LLC v. VeriFone, Inc.*, 807 F.3d 1346, 1350 (Fed. Cir. 2015). The district court construed parsing as "an automated process of analyzing a string according to a set of rules of a grammar."

A40. The court refused, however, to instruct the jury on the boundaries of "parsing" in Sprint's proposed construction that are compelled by the prosecution history: "Parsing does not include mapping a telephone number to a domain name, nor a lookup from a telephone number to a domain name, and it does not refer to any predetermined association between a number string and a domain name. Parsing is also not the same as translating." A38. Because the jury was not asked to consider whether Sprint's accused call flows practiced anything more than mapping, lookup, predetermined association, or translation, remand is necessary to hold Comcast to its burden.

A. The Prosecution History Reflects The Ordinary Meaning Of Claim Terms And Reflects Meanings Disclaimed By The Applicant

When determining the meaning of claim terms, the prosecution history is relevant in two significant respects. First, the prosecution history forms an essential part of the intrinsic record that determines the ordinary meaning of those terms to one of ordinary skill in the art. The ordinary meaning of a term cannot be evaluated in a vacuum, but “must be considered in the context of all the intrinsic evidence, including the claims, specification, and prosecution history.” *Biogen Idec, Inc. v. GlaxoSmithKline LLC*, 713 F.3d 1090, 1094 (Fed. Cir. 2013).

The prosecution history is also relevant to the doctrine of prosecution disclaimer, which “preclude[s] patentees from recapturing through claim interpretation specific meanings disclaimed during prosecution.” *Omega Eng’g, Inc. v. Raytek Corp.*, 334 F.3d 1314, 1323 (Fed. Cir. 2003); *Vizio, Inc. v. Int’l Trade Comm’n*, 605 F.3d 1330, 1339 n.6 (Fed. Cir. 2010) (“[T]he prosecution history (or file wrapper) limits the interpretation of claims so as to exclude any interpretation that may have been disclaimed or disavowed during prosecution in order to obtain claim allowance.”). Prosecution disclaimer “promotes the public notice function of the intrinsic evidence and protects the public’s reliance on definitive statements made during prosecution.” *Omega Eng’g*, 334 F.3d at 1324.

Disclaimer attaches where “the alleged disavowing actions or statements made during prosecution [are] both clear and unmistakable.” *Id.* at 1326.

B. Any Proper Construction Of “Parsing” Must Distinguish “Predetermined Associations,” “Mapping,” “Lookup,” And “Translating” Because The File History Expressly Distinguishes Those Processes When Defining “Parsing”

During reexamination of the ’916 and ’046 patents, as well as related patents by the same inventor, Comcast told the examiner what “parsing” is and what “parsing” is not. Those statements bind Comcast here. *See Ormco Corp. v. Align Tech., Inc.*, 498 F.3d 1307, 1314 (Fed. Cir. 2007) (relying on the doctrine of prosecution disclaimer based on statements in the prosecution history of familial patents). In particular, to avoid rejection of the claims, Comcast distinguished “parsing” from “predetermined association” and other processes such as “mapping,” “lookup,” and “translating.” Those distinctions should have been incorporated in the district court’s instructions to the jury, thereby binding Comcast to its clear statements establishing the scope of parsing. *See RFID Tracker, Ltd. v. Wal-Mart Stores, Inc.*, 342 F. App’x 628, 632 (Fed. Cir. 2009) (affirming district court’s reliance on statement in the prosecution history to construe the term “interrogator/reader” as meaning “an interrogator/reader that includes a field generator and a receiver, but not a transmitter”); *N. Am. Container, Inc. v. Plastipak Packaging, Inc.*, 415 F.3d 1335, 1343, 1346 (Fed. Cir. 2005)

(defining “generally convex” to require “a majority of convex points along the inner wall and *no* concave points” based on statements in prosecution history); *Southwall Techs., Inc. v. Cardinal IG Co.*, 54 F.3d 1570, 1574 (Fed. Cir. 1995) (defining “directly contiguous” using the negative limitations proffered by the patentee during prosecution).

During reexamination, Comcast made clear that the ordinary meaning of “parsing” does not refer to any predetermined association between a number string and a domain name. The examiner initially rejected the asserted claims of the ’916 and ’046 patents during reexamination. In response, Comcast submitted the Declaration of Dr. Zygmunt Haas, who disputed the examiner’s broad understanding of the term and clarified the proper definition according to Comcast in a section titled “Meaning of the Term ‘Parsing’”:

In reviewing the Office action rejecting these claims based on RFC 1703 and Schwartz I and II, it appears that the Office has interpreted “parsing” a number string to refer to any predetermined association between the number string and a domain name. Such an interpretation is inconsistent with how one of ordinary skill in the art at the time of the earliest priority date in 1996, having read the patent specification and prosecution history, would have understood the term “parsing.”

A2310; *see also* A2531, 2591, 2665, 2746-47. The meaning of “parsing” in the ’916 and ’046 patents “must be considered in the context of” that prosecution history. *Biogen Idec*, 713 F.3d at 1094. Either Comcast expressly defined “parsing” during prosecution or, alternatively, Dr. Haas’ statement constitutes a

clear and unmistakable disclaimer that “parsing” does not refer to any predetermined association between a number string and a domain name. *See Ecolab, Inc. v. FMC Corp.*, 569 F.3d 1335, 1342 (Fed. Cir. 2009) (“[S]ince, by distinguishing the claimed invention over the prior art, an applicant is indicating what the claims do not cover, he is by implication surrendering such protection.”). In either event, the construction of parsing should reflect that definitional limitation.

During reexamination, Comcast similarly distinguished the meaning of “parsing” from “mapping,” “lookup,” and “translation.” In response to a prior art rejection for the related ’304 patent, Dr. Haas rejected the Patent Office’s reading of “parsing” that included “mapping.” A2664.¹⁸ Likewise, Comcast distinguished prior art in the ’046 patent reexamination by arguing that “parsing” does not include a “lookup from a phone number to a domain name.” *See* A2959 (“[prior art] also does not provide parsing, as it is at best a mere lookup from a phone number to a domain name”). Comcast similarly overcame a double-patenting rejection by insisting that “[t]ranslating” is not the same as “parsing.” A2327,

¹⁸ *See also* A2939-40 (“Mapping, as already explained above, is not the same as parsing.”); A2893-94 (same); A2868 (distinguishing a prior art reference because it “lacks parsing and merely maps a telephone number to a domain name like a lookup table”).

3003.¹⁹ Comcast’s excision of the terms “mapping,” “lookup,” and “translation” from the meaning of “parsing” is definitional, or at least a clear and unequivocal disclaimer of claim scope that Comcast cannot recapture. *See, e.g., Omega Eng’g*, 334 F.3d at 1327 (holding that distinctions over prior art constituted a “deliberate surrender of claim scope, unmistakable in its effect because it is not suitable to multiple interpretations”).

C. The District Court Failed To Give Effect To The Prosecution History

Despite those clear and unmistakable statements in the prosecution history, the district court refused to instruct the jury on the limits of the term “parsing” in the ’916 and ’046 patents. That was error.

First, the district court doubted “that use of those negative limitations rose to the level of a clear disavowal of any claim scope contained within the plain meaning of parsing.” A40. The court did not explain the basis for its hesitation or point to any unclear disavowals. Even if the negative limitations were only relevant to prosecution disclaimer, Comcast’s disavowals were clear and unmistakable for the reasons described above, and this Court reviews prosecution

¹⁹ *See also* A2327 (“Parsing is a process of analyzing a string according to a set of rules. Conversely, “translating” generally refers to converting information from one form to another and does not necessarily imply any particular mechanism for that conversion.”).

disclaimer de novo. *Shire Dev., LLC v. Watson Pharm., Inc.*, 787 F.3d 1359, 1365 (Fed. Cir. 2015).

More importantly, the district court should have considered the prosecution history as evidence of the ordinary meaning of “parsing” as used in the ’916 and ’046 patents. Prosecution disclaimer requires a clear and unmistakable disavowal because a “heavy presumption exists that claim terms carry their full ordinary and customary meaning.” *Epistar Corp. v. Int’l Trade Comm’n*, 566 F.3d 1321, 1334 (Fed. Cir. 2009). But there is no reason to apply that standard when the prosecution history is evidence of the ordinary meaning itself, rather than a basis for *narrowing* the ordinary meaning for purposes of the asserted claims. *See Biogen Idec*, 713 F.3d at 1095 (“[W]hen the patentee unequivocally and unambiguously disavows a certain meaning to obtain a patent, the doctrine of prosecution history disclaimer narrows the meaning of the claim consistent with the scope of the claim surrendered.”).

The district court also refused to give effect to the disavowals in the intrinsic record because “the cited references are not in the record, and Sprint therefore has not shown any disclaimer.” A40. That was an unsupported non sequitur. The district court cited no case requiring that the record contain the cited references—as opposed to the disavowing statements—and Sprint is aware of none. To the contrary, this Court has held that prosecution disclaimer attaches to the applicant’s

statements in the prosecution history. *See Uship Intellectual Props., LLC v. United States*, 714 F.3d 1311, 1315 (Fed. Cir. 2013) (“Our cases broadly state that an applicant’s statements to the PTO characterizing its invention may give rise to a prosecution disclaimer.”). Indeed, the disavowal analysis “focuses on *what the applicant said*, not on whether the representation was necessary or persuasive.” *Id.* (emphasis added). Even if the cited references were relevant to the motives of the applicant or examiner (though the district court did not conclude that they were), “arguments made during prosecution shed light on *what the applicant meant by its various terms*.” *Id.* at 1316 (emphasis added). Because the record already contains the applicant’s “disavowing actions or statements,” *Omega Eng’g*, 334 F.3d at 1325-26, nothing more is necessary to find that Comcast disavowed mapping, lookup, and predetermined associations in the definition of parsing.

IV. The District Court Awarded Windfall Prejudgment Interest During A Time When Two Patents Could Not Have Been Infringed Because They Did Not Exist

A. Prejudgment Interest Is Available To Compensate During Periods Of Infringement

A district court has discretion to award prejudgment interest to compensate the patent owner for infringement. A party that has proven patent infringement is authorized by statute to recover “damages adequate to *compensate* for the infringement . . . together with interest.” 35 U.S.C. § 284 (emphasis added).

While damages compensate the patent owner for the royalty payments he would have earned, a court awards prejudgment interest to compensate the patent owner for the “forgone use of the [payments] between the time of infringement and the date of the judgment.” *Gen. Motors Corp. v. Devex Corp.*, 461 U.S. 648, 656 (1983).

The compensatory purpose of prejudgment interest requires that any award be “from the date of infringement to the date of judgment.” *Nickson Indus., Inc. v. Rol Mfg. Co.*, 847 F.2d 795, 800 (Fed. Cir. 1988); *see also Oiness v. Walgreen Co.*, 88 F.3d 1025, 1033 (Fed. Cir. 1996) (“Interest compensates the patent owner for the use of its money between the date of injury and the date of judgment.”). The district court must tether prejudgment interest to the date infringement began because there is no need for compensation until the patent has been infringed. That is true whether damages are based on a hypothetical royalty negotiation or some other measure. *LaserDynamics, Inc. v. Quanta Computer, Inc.*, 694 F.3d 51, 75 (Fed. Cir. 2012) (“In general, the date of the hypothetical negotiation is the date that the infringement began.”).

Moreover, a patent owner is not entitled to prejudgment interest for the period of time before the infringed patent was issued. “[A] party cannot be held liable for ‘infringement’ . . . of a *nonexistent* patent, *i.e.*, no damages are payable on products manufactured and sold *before the patent issued.*” *Gustafson, Inc. v.*

Intersystems Indus. Prods., Inc., 897 F.2d 508, 510 (Fed. Cir. 1990) (last emphasis added). Since a patent cannot be infringed until it has been issued, a district court should not compensate a patent owner for infringement via prejudgment interest until the patent has been issued.

B. The District Court's Award Overcompensates Comcast By Granting Interest Before There Could Have Been Infringement

The district court erroneously awarded prejudgment interest before there could have been infringement for two of the patents-in-suit. The '916 patent was issued in 2006, but both the '008 and the '046 patents were issued in 2012. Despite the different issue dates, the district court awarded prejudgment interest for *all* three patents starting in 2006. A65-66. That award overcompensated Comcast because Sprint could not have legally infringed the '008 or '046 patents until they were issued in 2012. The district court's order violated the cardinal rule that prejudgment interest must be "based on a hypothetical negotiation at the time that infringement began, not an earlier one." *Applied Med. Res. Corp. v. U.S. Surgical Corp.*, 435 F.3d 1356, 1361 (Fed. Cir. 2006). This error was particularly significant because the '008 patent represented more than 90% of the proposed damages award calculated by Comcast's expert. A6560. Thus, nearly all of the prejudgment interest accrued before Sprint could have infringed the most valuable patent in this litigation.

The district court misapplied the principle that when using a hypothetical negotiation to calculate damages, “the district court may . . . consider the panoply of events and facts that occurred thereafter and that could not have been known to or predicted by the hypothesized negotiators.” A66 (quoting *ResQNet.com, Inc. v. Lansa Inc.*, 594 F.3d 860, 872 (Fed. Cir. 2010)). The district court thought that “[a] hypothetical negotiation in 2006 . . . would include the fact that two of the patents [the ’008 and the ’046] issued later.” A66. In other words, the district court believed that the 2012 patents are subsequent “facts” that the hypothetical negotiators could consider in 2006.

That new patents will be awarded in the future is not the sort of “fact” that justifies the premature accumulation of interest. The principle cited by the district court originated in *Fromson v. W. Litho Plate & Supply Co.*, 853 F.2d 1568 (Fed. Cir. 1988), *overruled on other grounds by Knorr-Bremse Sys. v. Dana Corp.*, 383 F.3d 1337 (Fed. Cir. 2004) (en banc). *Fromson* cautioned that relying on future events that arose after the hypothetical negotiation date must be done with care. The court held that “a court is not at liberty, in conducting the [hypothetical negotiation] methodology, to abandon entirely the statutory standard of damages ‘adequate to compensate’ for the infringement.” *Fromson*, 853 F.2d at 1575. It emphasized that a reasonable royalty based on a hypothetical negotiation “must be at least a close approximation of what would be ‘adequate to compensate’ for the

‘use made of the invention by the infringer.’” *Id.* (quoting 35 U.S.C. § 284). Under *Fromson*, then, the “facts” to be considered that arose after the hypothetical negotiation date must be tethered to the statutory purpose of providing compensation to the patent owner. The key question under this approach is whether taking notice of subsequent facts is necessary to produce “a close approximation of what would be ‘adequate to compensate’” the patentee. *Id.* Here, because Comcast had *no right to enforce* the ’008 and ’046 patents in 2006, that those patents would be issued in the future is irrelevant to what would be adequate to compensate Comcast before 2012.

Sinclair Refining Co. v. Jenkins Petroleum Process Co., 289 U.S. 689 (1933), also undermines the district court’s decision. In *Sinclair*, the Court noted that events and facts arising after the hypothetical negotiation date are “a book of wisdom that courts may not neglect.” 289 U.S. at 698. But the Court emphasized that “[t]o correct uncertain prophecies in such circumstances is not to charge the offender with elements of value nonexistent at the time of his offense.” *Id.* That, however, is *exactly* what the district court did to Sprint. It awarded prejudgment interest starting in 2006 for patents that did not issue until 2012. This amounted to charging Sprint for “elements of value nonexistent at the time of [its] offense.” *Id.* When the parties in this case hypothetically sat down to negotiate in 2006, any value from the 2012 patents had not yet materialized.

Tellingly, in the proceedings below, neither the district court nor Comcast cited a single case that used the *Sinclair-Fromson* principle as the district court did here. That is unsurprising: awarding prejudgment interest for infringing a patent before it has issued would systematically overcompensate patentholders. In a context very similar to the one here, this Court reaffirmed that principle in *DDR Holdings, LLC v. Hotels.com, L.P.*, 773 F.3d 1245 (Fed. Cir. 2014). There, the parties disputed whether the patent owner was entitled to prejudgment interest. *Id.* at 1262. As in this case, the district court awarded prejudgment interest on multiple patents despite the fact that one of them was issued in 2010 and the other was issued in 2006. *Id.* at 1263. On appeal, this Court held that the patent which issued in 2006 was invalid, and that “the district court must recalculate its award of prejudgment interest so that it is tied solely to [the defendant’s] infringement of the ‘399 patent, which issued in 2010, more than four years after issuance of the ‘572 patent.” *Id.*

This Court recognized in *DDR* that prejudgment interest cannot be awarded during a period of time when a patent could not have been legally infringed. After this Court held that one of the patents there was invalid, it directed the district court to recalculate its prejudgment interest award. Just as one cannot infringe an invalid patent, one cannot infringe an unissued patent. *Gustafson, Inc.*, 897 F.2d at 510. Here, the district court should have adjusted the prejudgment interest award

to account for the fact that two of the patents were not issued until 2012. Its failure to do so rewarded Comcast for “infringement” that it could not legally suffer. Accordingly, the district court’s prejudgment interest award must be reversed because it is inconsistent with *DDR* and fundamental principles of patent law.

C. There Is Ample Foundation In This Record To Calculate Prejudgment Interest Properly

The district court thought it was powerless to determine what portion of the jury’s award was attributable to each patent. A66. But Comcast’s own expert provided the district court with a proposed damages award broken down by patent. A6560. Although the jury ultimately awarded about 45 percent of what Comcast’s expert proposed, it is a matter of simple arithmetic to prorate the amount proposed for each patent by the jury’s award. Sprint provided that arithmetic to the court. A6560.

The district court was fully empowered to use the prorated proposed damages from Comcast’s expert for purposes of prejudgment interest. While damages are determined by the jury, prejudgment interest is within the court’s broad discretion. By concluding that it lacked authority to use the calculation implied by Comcast’s expert, the court committed legal error that constituted an abuse of discretion. *See Lummus Indus., Inc. v. D.M. & E. Corp.*, 862 F.2d 267,

274 (Fed. Cir. 1988) (rejecting district court's reason for not awarding prejudgment interest because it was "legally erroneous").

CONCLUSION

The jury verdict must be vacated because it is based on legally improper theories of infringement with respect to the '008, '916, and '046 patents. The case should be remanded for further proceedings. If this Court determines that the jury verdict is proper in part and improper in part, then those further proceedings must include an allocation of damages among the asserted patents. Finally, the case should also be remanded for a recalculation of prejudgment interest.

Dated: April 4, 2016

Respectfully Submitted,

/s/ Brian D. Schmalzbach

Brian C. Riopelle
David E. Finkelson
Brian D. Schmalzbach
MCGUIREWOODS LLP
800 East Canal Street
Richmond, VA 23219
Tel: (804) 775-1084

Rachelle H. Thompson
MCGUIREWOODS LLP
434 Fayetteville Street
Raleigh, NC 27601
Tel: (919) 755-6572

*Counsel for Sprint Communications
Company L.P., Sprint Spectrum L.P., and
Nextel Operations, Inc.*

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

COMCAST IP HOLDINGS I, LLC, :
 :
 Plaintiff, :
 :
 v. : C.A. No. 12-205-RGA
 :
 SPRINT COMMUNICATIONS COMPANY :
 L.P.; SPRINT SPECTRUM L.P.; and :
 NEXTEL OPERATIONS, INC., :
 :
 Defendants. :

CLAIM CONSTRUCTION

Arthur G. Connolly, Esq., Wilmington, Delaware; David J. Lisson, Esq. (argued), Menlo Park, California; Anthony I. Fenwick, Esq. (argued), Menlo Park, California; Attorneys for Plaintiff Comcast IP Holdings I, LLC Technology, Inc.

Richard K. Herrmann, Esq., Wilmington, Delaware; Robert H. Reckers, Esq. (argued), Houston, Texas; Peter E. Strand, Esq. (argued), Washington, D.C.; Attorneys for Defendants Sprint Communications Company L.P., Sprint Spectrum L.P., and Nextel Operations, Inc.

August 16, 2013
Wilmington, Delaware


ANDREWS, UNITED STATES DISTRICT JUDGE:

This is a claim construction opinion. Plaintiff Comcast IP Holdings I, LLC (“Comcast”) asserts certain patent rights against Defendants Sprint Communications Company L.P., Sprint Spectrum L.P., and Nextel Operations, Inc. (collectively, “Sprint”).¹ Those rights include one group of patents, referred to by the parties as the “Low Patents,” and an unrelated patent, U.S. Patent No. 6,873,694. The “Low Patents” are U.S. Patent Nos. 7,012,916, 7,206,304, 7,903,641, 8,170,008, 8,189,565, 8,204,046, and 8,223,752.

The “Low Patents” claim inventions facilitating the integration of traditional telephone networks with computer networks. The ‘694 Patent claims the invention of systems and methods to optimize a telephony network by accounting for different telephony parameters to determine whether a request for telephony network service should be accepted.

A. “Low Patents”

1. “switched telecommunication system / telecommunication system”

¹ Defendant Sprint Communications in turn counterclaimed that Comcast infringes six of its patents. That litigation is proceeding separately in Case No. 12-1013.

Claim term	Comcast's Construction	Sprint's Construction
"switched telecommunication system" 8,170,008: Claims 1, 3, 6, 7, 12, 13, 19, 27, and 29; 8,189,565: Claims 1 and 3; and 8,223,752: Claims 16, 24, and 30.	A system comprising a bearer network with switches for setting up a bearer channel through the network. A datagram-based communication system where each data packet is independently routed through a bearer network without following a predetermined bearer channel is not a "switched telecommunication system."	A system comprising a bearer network with switches for setting up a bearer channel through the network that does not include datagram-based communication systems where each data packet is independently routed through a bearer network without following a predetermined bearer channel.
"telecommunication system" 8,170,008: Claims 19 and 29.	A system comprising a bearer network with switches for setting up a bearer channel through the network. A datagram-based communication system where each data packet is independently routed through a bearer network without following a predetermined bearer channel is not a "telecommunication system."	A system comprising a bearer network with switches for setting up a bearer channel through the network that does not include datagram-based communication systems where each data packet is independently routed through a bearer network without following a predetermined bearer channel.

The parties dispute the construction of "switched telecommunication system." They do agree that a "switched telecommunication system" is defined in the specification as "a system comprising a bearer network with switches for setting up a bearer channel through the network." See '008 Patent at 1:31-33. They differ over whether a "switched telecommunication system" may include certain additional elements, those being elements of a "datagram-based communication system[]" where each data packet is independently routed through a bearer network without following a predetermined bearer channel" ("datagram-based system"). *Id.* at 1:52-55. Sprint argues that the specification specifically defines the "switched telecommunication system" to exclude the elements of a "datagram-based system." This is

because the specification states that a “communication system” in general is a broader concept than the “switched based telecommunication system” because a “communication system” may include a “datagram-based system.” *See id.* at 1:50-55. Comcast disagrees, arguing that a “communication system” is broader than a “switched telecommunication system” because a “communication system” includes a system without switches, but this does not necessarily mean that a “switched telecommunication system” cannot also have aspects of a “datagram-based system.” That is, so long as a “switched telecommunication system” indeed has switches, it may also have aspects of a “datagram-based system.”

The Court agrees with Comcast. A system with elements of both switches and a “datagram-based system” is not necessarily outside the scope of a “switched telecommunication system.” The fact that the “communication system” is understood to be broader than the “switched telecommunication system” does not mean a “switched telecommunication system” is precluded from having some overlap with elements of a “datagram-based system.” A “communication system” is still a broader concept than a “switched telecommunication system” even where the “switched telecommunication system” has elements of a “datagram-based system,” because a “switched telecommunication system” at a minimum must have switches and function on a bearer network. A “communication system” does not have those requirements. Sprint’s argument that a system would no longer be a “switched telecommunication system” so long as it includes datagram elements, even if it indisputably includes switches and is on a bearer network, is not justified by the claims or specification.

2. “requesting. . .a communication to be set up through the switched telecommunication system”

Claim term	Comcast’s Construction	Sprint’s Construction
------------	------------------------	-----------------------

“requesting . . . a communication to be set up through the switched telecommunication system” 8,189,565: Claim 1	Requesting a communication to be set up through a bearer network or a signaling network of the switched telecommunication system.	Requesting a communication to be set up through a bearer channel of the switched telecommunication system.
---	---	--

At oral argument, the parties agreed to the plain and ordinary meaning of this term. (D.I. 97, p. 48).

3. “telecommunications system control apparatus”

Claim term	Comcast’s Construction	Sprint’s Construction
“telecommunications system control apparatus” 7,903,641: Claims 13, 17, 21, and 25.	A device involved in the processing of signaling used in a telecommunications system to effect call control.	An apparatus that controls a “telecommunication system”

At oral argument, the parties agreed to the following construction of this term: “a device that processes signaling used in a telecommunications system to effect call control.” (D.I. 97, p. 34).

4. “URI / uniform resource identifier (URI) / universal resource identifier (URI) / universe resource name (URN)”

Claim term	Comcast’s Construction	Sprint’s Construction
“URI / uniform resource identifier (URI) / universal resource identifier (URI) / universe resource name (URN)” 7,012,916: Claim 45;	No construction necessary. In the alternative: “A compact string of characters used to identify a resource accessible over a network.”	An internet-resolvable indicator of a location of a phone page.

8,204,046: Claims 90 and 113; 7,903,641; Claims 13 and 21	Modified alternative: “A compact string of characters used to identify a resource accessible over a network and adhering to a syntax in which a naming scheme specifier is followed by a string whose format is a function of the naming scheme, with the name of the scheme separated from the following string by a colon.	
“uniform resource name” 8,170,008: Claim 28.	No construction necessary. In the alternative: “A uniform resource identifier (URI) that identifies the resource by name”	A Uniform Resource Identifier (URI) (as that term is construed by the Court) that identifies the resource by name.

The parties agree that “URI” is a term of art with a well-understood meaning. They dispute whether the meaning is redefined by the specification. Sprint argues that the following statement in the “Best Mode” section of the patent redefined the scope of “URI:” “(for convenience, the more general term URI will be used hereinafter to mean the Internet-resolvable indicator of a phone page)”. ’916 Patent at 13:22-25. Comcast argues that this narrower definition only applies to the “Best Mode” section and should not limit the term as it is used generally in the claims of the patent, especially because “URI” is used in sections prior to the “Best Mode” section with no redefinition or limitation.

The Court agrees with Comcast. “URI” is a widely used term with a long history in the art of computer networks, and the Court sees no clear sign that the patentee intended a redefinition. First, the passage pointed to by Sprint styles itself to be only used “for convenience,” suggesting it is not a permanent redefinition and lessening the persuasiveness of its complete applicability to the claims of the patent. The “for convenience” qualifier suggests a

desire to simplify the discussion in the particular context of the “Best Mode” section, which is an extremely detailed discussion of a preferred embodiment.² *See id.* at 12:52-53. The understanding of the patentee’s intentions is supported by his willingness to introduce the term “URI” with its traditional meaning earlier in the specification. *See, e.g., id.* at 7:61-66.

The inference that this is not a definition is bolstered by its comparison with truly clear cut definitions of the specification. Those definitions are found in the “Field of the Invention” section at the very start of the specification, fall within quotation marks, and are accompanied by the words “when used herein.” Sprint’s proposed redefinition shares none of those qualities. Finally, independent claim 90 of the ‘046 Patent includes the term “uniform resource identifier (URI),” and corresponding dependent claim 103 limits “URI” in just one way, to “an Internet-resolvable indicator of a location of a phone page.” The doctrine of claim differentiation suggests that “URI” should be construed so that the construction results in independent claim 90 being broader than dependent claim 103. The only way to do that is if “URI” is not limited to just “an internet resolvable indicator of a location of a phone page.”

Sprint argues that the “hereinafter” language in the alleged redefinition of “URI” requires the Court to apply the redefinition in all places subsequent, including the claims. While this is not an unreasonable stance if the phrase existed in a vacuum, or if the specification were devoid of other references to “URIs,” the Court finds that it would be misplaced to find a clear definition when viewed in the context of the entire specification. As to the prosecution history cited by

² By my count, the term “URI” is used at least 65 times in the “Best Mode” section.

Sprint, it represents statements of the examiner rather than the patentee, and thus cannot limit the claims. The Court thus adopts the plain and ordinary meaning of “URI.”³

5. “DNS-type database system / DNS-type distributed database system”

Claim term	Comcast’s Construction	Sprint’s Construction
<p>“a DNS-type database system”</p> <p>7,012,916: Claim 45</p>	<p>A database system comprised of one or more servers that associates domain names with one or more records and uses an IP protocol and a pre-determined message format, wherein the domain names are hierarchically structured and at least one server of the system may be addressed through a resolver.</p>	<p>A system having the following characteristics of the Domain Name System:</p> <p>i) host name space is organized as a tree structured hierarchy of nodes with each host having a corresponding leaf node; each node has a label (except the root node) and each label begins with an alphabetic character and is followed by a sequence of alphabetic characters or digits;</p> <p>ii) each host has one or more associated Registration Records (“RR”);</p> <p>iii) There are one or more DNS servers each with responsibility for a subtree of the name space. A DNS server will hold RRs for all or part of its subtree--in the latter case it delegates responsibility for the remainder of the subtree to one or more further DNS servers. A DNS server knows the address of any server to which it has delegated responsibility and also the address of the server which has given it the responsibility for the subtree it manages. The DNS servers thus point to each other in a structuring reflecting that of the naming hierarchy;</p> <p>iv) An application wishing to make use of the DNS does so through an</p>

³ Should the parties disagree as to the plain and ordinary meaning of this term, they will be given additional opportunity to argue it.

		<p>associated “resolver” that knows the address of at least one DNS server. When a DNS server is asked by this resolver for an RR of a specified host, it will return either the requested RR or the address of a DNS server closer to the server holding the RR in terms of traversal of the naming hierarchy. In effect, the hierarchy of the servers is ascended until a server is reached that also has responsibility for the domain name to be resolved; thereafter, the DNS server hierarchy is descended down to the server holding the RR for the domain name to be resolved.</p> <p>v) using a predetermined message format and IP protocols.</p>
<p>“a DNS-type distributed database system”</p> <p>7,206,304: Claim 7.</p>	<p>A database system comprised of a plurality of servers that associates domain names with one or more records and uses an IP protocol and a pre-determined message format, wherein the domain names are hierarchically structured and at least one server of the system may be addressed through a resolver.</p>	<p>A system having the following characteristics of the Domain Name System:</p> <p>i) host name space is organized as a tree structured hierarchy of nodes with each host having a corresponding leaf node; each node has a label (except the root node) and each label begins with an alphabetic character and is followed by a sequence of alphabetic characters or digits;</p> <p>ii) each host has one or more associated Registration Records (“RR”);</p> <p>iii) There a plurality of DNS servers each with responsibility for a subtree of the name space. A DNS server will hold RRs for all or part of its subtree--in the latter case it delegates responsibility for the remainder of the subtree to one or more further DNS servers. A DNS server knows the address of any server to which it has delegated responsibility and also the address of the server which has given it the</p>

		<p>responsibility for the subtree it manages. The DNS servers thus point to each other in a structuring reflecting that of the naming hierarchy;</p> <p>iv) An application wishing to make use of the DNS does so through an associated “resolver” that knows the address of at least one DNS server. When a DNS server is asked by this resolver for an RR of a specified host, it will return either the requested RR or the address of a DNS server closer to the server holding the RR in terms of traversal of the naming hierarchy. In effect, the hierarchy of the servers is ascended until a server is reached that also has responsibility for the domain name to be resolved; thereafter, the DNS server hierarchy is descended down to the server holding the RR for the domain name to be resolved.</p> <p>v) using a predetermined message format and IP protocols.</p>
--	--	---

The next term is “a DNS-type database system/ DNS-type distributed database system.” Sprint argues that the term has been clearly defined by the specification, as written from 5:47-6:24 of the '916 Patent and copied above. The specification states that the characteristics of Sprint’s proposed construction represent the “main characteristics of the DNS” and also that they “may be considered as defining a ‘DNS-type’ system always allowing for minor variations such as in label syntax, how the labels are combined (ordering, separators), the message format details, evolution of the IP protocols etc.” *Id.* at 5:45-47, 6:25-29. Sprint argues that its proposal is an explicit definition and should be adopted in its entirety. Comcast argues that Sprint’s construction is prohibitively long, and the Court should instead adopt its version, which is truncated but nevertheless provides a concise summary that includes each element of a DNS-type

system understandable to a jury. Comcast further argues that Sprint's construction fails to account for the statement that "minor variations" from that proposed definition are permissible. In response, Sprint argues that Comcast's construction is incomplete, inexplicably leaving out numerous "main characteristics of the DNS," including the inventors' reference to a tree-structured hierarchy, the requirement of "Registration Records," and how delegation and addressing are handled. Sprint also argues that the doctrine of equivalents analysis accounts for the "minor variations" language in the specification.

The Court agrees with Sprint. Comcast may now regret the patentee's decision to explicitly define the claim term with a long and unwieldy definition, but that is the consequence with which it must live. The patentee used definitional language in connection with this term, by explicitly stating that the long quotation "may be considered as defining a 'DNS-type' system." Comcast does not cite any intrinsic evidence explaining why its proposed definition encompasses the key elements of a DNS-type system, nor does Comcast cite evidence explaining why omissions are mere "minor variations" that need not be incorporated into the construction. As pointed out by Sprint, "minor variations" from the claim may be accounted for by the doctrine of equivalents.

6. Domain name system signaling

Claim term	Comcast's Construction	Sprint's Construction
"domain name system signaling" 8,170,008: Claims 1, 5, 8, 13, 16, 27 and 28.	Signaling exchanged with the Domain Name System ("DNS") of the Internet or with a DNS-type database system.	A message format of the Domain Name System. Modified proposal: "a DNS-formatted message of the Domain Name System."

The next term is “domain name system signaling.” The first dispute, whether the term is limited to the Domain Name System of the Internet, or whether it also includes DNS-type database systems, is now resolved, as the parties agreed that the term encompasses both systems. (D.I. 97, p. 82). The remaining dispute thus boils down to whether the “domain name system signaling” encompasses any signal exchanged on the system, as argued by Comcast, or only a message formatted by the DNS, as argued by Sprint.

Sprint argues that its construction should be adopted because every description in the patent about DNS or “DNS-type systems” involves a message format for querying the DNS database and receiving a response. ‘008 Patent at 5:43-56; Figures 3-5, 6:49-7:18. Sprint also points out that one of the five defined “main characteristics” of the “DNS-type system” includes its use of a “predetermined message format.” *Id.* at 5:56-58. Of the allowable “minor variations” of the “DNS-type system,” one includes the “message format details.” *Id.* at 5:63. Comcast concedes that DNS-type systems use “a predetermined message format, and also concedes that signaling and messaging are synonymous. (D.I. 81, p. 75). Comcast disagrees, however, that the “DNS-type system” itself formats the message. There is, however, no suggestion within the specification as to which other component of the invention would format the message. It is not disputed that the “DNS-type system” is the component that sends and receives the messages, and that the system’s use of a predetermined message format is a main characteristic of the system. It is further not disputed that the “message format details” are varied within the “DNS-type system” itself, not some other component of the invention. It follows that, as the “message format” is controlled by the “DNS-type system,” the “DNS-type system” formats the actual message. For these reasons, the Court construes “domain name

system signaling” as “a DNS-formatted message of the Domain Name System of the Internet or with a DNS-type system.”

7. A substantial portion of the number string

Claim term	Comcast’s Construction	Sprint’s Construction
“a substantial portion of the number string” 7,012,916: Claim 45; and 7,206,304: Claim 7.	A portion of a number string having a distinct meaning such as, in the case of a telephone number, the country code, the area code, or the local number.	Indefinite.

The next term is “a substantial portion of the number string.” The dispute as to this term is whether it can be construed at all, as Sprint argues that it is indefinite. Sprint argues that a “number string” is generically claimed with an unlimited scope that can be practiced by an infinite number of strings with various lengths, groupings, and meanings, including phone numbers, local routing numbers, groupings, meanings, etc. In response, Comcast argues that the claims themselves offer context to understand the scope of the “number string,” as they are associated with identifying a target entity. *See, e.g.*, ‘916 Patent at claim 45; ‘304 Patent at claim 7. Comcast further argues that a person skilled in the art would understand the “distinct meaning” any individual number string may have according to the distinct groups of a number. In response, Sprint argues that a “number string” is not fairly limited to the only type of string described as parsed in the patents.

The Court agrees with Comcast. Although the term “number string” is vague in a vacuum, it must be read in light of the specification as a whole. The only type of strings described in the patents are telephone numbers, and this helps inform the construction of the term. The Court

thus does not believe that “number string” is so “insolubly ambiguous” as to evade construction and adopts Comcast’s proposed construction. The Court is not making an ultimate finding as to the indefiniteness of the term, and Sprint is free to renew the argument at the summary judgment stage.

B. U.S. Patent No. 6,873,694

The remaining three terms are from the claims of the ‘694 Patent.

1. Dial-up prompt [parameter]

Claim term	Comcast’s Construction	Sprint’s Construction
“dial-up prompt [parameter]” Claim 5	A parameter that allows a user to decide whether to allow network connections on each request or whether a prompt to the user is required on each request.	A parameter that allows a requesting application, or requester, to give a user of an appliance control over an Internet connection.

The parties dispute the construction of “dial-up prompt [parameter].” The key dispute in scope is whether the “dial-up prompt [parameter]” must trigger a prompt in response to each attempt to connect to the network, as argued by Sprint, or whether it includes the ability to decide whether the prompt is required at all, as argued by Comcast.

The specification states, “A dial-up prompt parameter may allow a user to decide whether to allow network connections on each request; that is, the user may decide when it is appropriate to connect to the telephony network using an appliance, and when to keep the telephone available for phone conversations.” ‘694 Patent at 3:17-22. The specification is clear that the “dial up prompt parameter” arms the user with the ability to “decide when it is appropriate to connect to the telephony network.” This is a broad-sounding statement that, most naturally read,

would include any decision controlling connectivity to the network, even one made in advance of a request that did not require the presentation of a prompt. It does not require the decision to allow or deny a network connection to be made in response to or contemporaneously with each instance of request. This is supported by Figure 2 of the specification, which shows that the user may control whether a prompt is required.

Sprint's arguments are not convincing. Sprint relies on the following from the specification: "a dial-up prompting parameter, which allows a requesting application, or requester, to give a user of an appliance control over an Internet connection." *Id.* at 2:56-69. This sentence, however, merely states what the dial-up prompting parameter allows, and is not inconsistent with allowing advanced control or setting the prompt to not appear at all. Further, the claim does not require a "dial-up prompt," it requires a "dial-up prompt parameter," and the "parameter" may be used to control whether the prompt appears at all. Finally, Sprint argues that the Court's construction will incorrectly "ensnare systems that do not actually require dial-up prompts." This is not so. Any accused system must have the capability of presenting a dial-up prompt.

2. Telephony parameter

Claim term	Comcast's Construction	Sprint's Construction
"telephony parameter" Claims 1, 5, and 21.	A variable or factor used to determine whether to allocate a channel on a telephony network to an application or to a phone conversation.	Plain and ordinary meaning.

The parties next dispute the scope of "telephony parameter." Sprint argues for the plain meaning of the term, while Comcast argues that the term is not so simple as to justify a plain and

ordinary construction. The Court agrees with Sprint, as Comcast's construction only repeats requirements already found in two of three claims. For example, claim 1 states, "determining whether to allocate to the application a channel on the telephony network based on telephony parameters for obtaining balanced network service between the application usage of the telephony network and telephone usage of the telephony network." Where those requirements are not found, in claim 21, the Court agrees with Sprint that the claim provides its own context sufficient to understand the "telephony parameters" at play. As there does not seem to be an actual dispute of scope here, *i.e.*, neither party genuinely represents that anything of significance is wrongly excluded or included by the other's construction, the Court adopts the plain and ordinary meaning of "telephony parameter."

3. Telephone usage of the telephony network

Claim term	Comcast's Construction	Sprint's Construction
"telephony usage of the telephony network"	Use of the telephony network by a telephone for a phone conversation	Plain and ordinary meaning
Claim 1		

The final term is "telephone usage of the telephony network." The dispute of scope here is whether such "telephone usage" is restricted to network usage for phone conversations, as argued by Comcast, or whether it may also include other phone activities, such as voicemail or text messages, as argued by Sprint. The Court agrees with Sprint. The use of a telephone on a telephony network most naturally includes more functionality than a mere voice conversation, even going back to the 2001 filing date of this patent, when voicemail and text messages were already well-known. The Court thus adopts the plain and ordinary meaning of this term.

The parties are instructed to jointly submit a claim construction order suitable for submission to the jury consistent with this opinion within 14 days.

**UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

COMCAST IP HOLDINGS I, LLC,

Plaintiff,

v.

SPRINT COMMUNICATIONS COMPANY
L.P.; SPRINT SPECTRUM L.P.; and
NEXTEL OPERATIONS, INC.;

Defendants.

Case No.: 1:12-cv-0205-RGA (CJB)

DEMAND FOR JURY TRIAL

CLAIM CONSTRUCTION ORDER

After having considered the submissions of the parties and hearing oral argument on the matter, IT IS HEREBY ORDERED, ADJUDGED, and DECREED that:

1. As used in the asserted claims of U.S. Patent Nos. 7,012,916, 7,206,304, 7,903,641, 8,170,008, 8,189,565, 8,204,046, and 8,223,752:

The term “switched telecommunication system” is construed to mean “a system comprising a bearer network with switches for setting up a bearer channel through the network. A datagram-based communication system where each data packet is independently routed through a bearer network without following a predetermined bearer channel is not a ‘switched telecommunication system.’”

The term “telecommunication system” is construed to mean “a system comprising a bearer network with switches for setting up a bearer channel through the network. A datagram-based communication system where each data packet is independently routed through a bearer network without following a predetermined bearer channel is not a ‘telecommunication system.’”

Pursuant to the parties’ agreement, the term “requesting. . . a communication to be set up

through the switched telecommunication system” should be given its plain and ordinary meaning.

Pursuant to the parties’ agreement, the term “telecommunications system control apparatus” should be construed to mean “a device that processes signaling used in a telecommunications system to effect call control.”

The term “URI / uniform resource identifier (URI) / universal resource identifier (URI) / universe resource name (URN)” shall be given its plain and ordinary meaning.

The term “a DNS-type database system” is construed to mean “a system having the following characteristics of the Domain Name System:

- i) host name space is organized as a tree structured hierarchy of nodes with each host having a corresponding leaf node; each node has a label (except the root node) and each label begins with an alphabetic character and is followed by a sequence of alphabetic characters or digits;
- ii) each host has one or more associated Registration Records (‘RR’);
- iii) there are one or more DNS servers each with responsibility for a subtree of the name space. A DNS server will hold RRs for all or part of its subtree—in the latter case it delegates responsibility for the remainder of the subtree to one or more further DNS servers. A DNS server knows the address of any server to which it has delegated responsibility and also the address of the server which has given it the responsibility for the subtree it manages. The DNS servers thus point to each other in a structuring reflecting that of the naming hierarchy;
- iv) an application wishing to make use of the DNS does so through an associated

'resolver' that knows the address of at least one DNS server. When a DNS server is asked by this resolver for an RR of a specified host, it will return either the requested RR or the address of a DNS server closer to the server holding the RR in terms of traversal of the naming hierarchy. In effect, the hierarchy of the servers is ascended until a server is reached that also has responsibility for the domain name to be resolved; thereafter the DNS server hierarchy is descended down to the server holding the RR for the domain name to be resolved;

- v) using a predetermined message format and IP protocols."

The term "a DNS-type distributed database system" is construed to mean "a system having the following characteristics of the Domain Name System:

- i) host name space is organized as a tree structured hierarchy of nodes with each host having a corresponding leaf node; each node has a label (except the root node) and each label begins with an alphabetic character and is followed by a sequence of alphabetic characters or digits;
- ii) each host has one or more associated Registration Records ('RR');
- iii) there are a plurality of DNS servers each with responsibility for a subtree of the name space. A DNS server will hold RRs for all or part of its subtree—in the latter case it delegates responsibility for the remainder of the subtree to one or more further DNS servers. A DNS server knows the address of any server to which it has delegated responsibility and also the address of the server which has given it the responsibility for the subtree it manages. The DNS servers thus point to each other in a structuring reflecting that of the

naming hierarchy;

- iv) an application wishing to make use of the DNS does so through an associated 'resolver' that knows the address of at least one DNS server. When a DNS server is asked by this resolver for an RR of a specified host, it will return either the requested RR or the address of a DNS server closer to the server holding the RR in terms of traversal of the naming hierarchy. In effect, the hierarchy of the servers is ascended until a server is reached that also has responsibility for the domain name to be resolved; thereafter the DNS server hierarchy is descended down to the server holding the RR for the domain name to be resolved;
- v) using a predetermined message format and IP protocols."

The term "domain name system signaling" is construed to mean "a DNS-formatted message of the Domain Name System of the Internet or with a DNS-type system."

The term "a substantial portion of the number string" is construed to mean "a portion of a number string having a distinct meaning such as, in the case of a telephone number, the country code, the area code, or the local number."

Pursuant to the parties' agreement, the term "setting up a call through the switched telecommunication system" should be construed to mean "setting up a bearer channel through a bearer network of the switched telecommunication system."

Pursuant to the parties' agreement, the term "initiating a call through the switched telecommunication system" should be construed to mean "initiating a communication through a bearer channel set up across a bearer network of the switched telecommunication system."

Pursuant to the parties' agreement, the term "setting up a call" should be construed to

mean “setting up a bearer channel through a bearer network of the switched telecommunication system.”

2. As used in the asserted claims of U.S. Patent No. 6,873,694:

The term “dial-up prompt [parameter]” is construed to mean “a parameter that allows a user to decide whether to allow network connections on each request or whether a prompt to the user is required on each request.”

The term “telephony parameter” shall be given its plain and ordinary meaning.

The term “telephone usage of the telephony network” shall be given its plain and ordinary meaning.

Pursuant to the parties’ agreement, the term “channel” should be construed to mean “a portion of bandwidth in a network that may be allocated to a request.”

Pursuant to the parties’ agreement, the term “priority parameter” should be construed to mean “a parameter indicating whether a request is immediate or deferrable.”

Pursuant to the parties’ agreement, the term “connection hold time parameter” should be construed to mean “a parameter defining how long a connection should be maintained after the last request has been released.”

Pursuant to the parties’ agreement, the term “request persistence parameter” should be construed to mean “a parameter defining whether a request may initiate a connection if one has not already been established, and whether multiple attempts are to be made if a connection attempt fails.”

Pursuant to the parties’ agreement, the term “application timeout parameter” should be construed to mean “a parameter used to define how long a requesting application can maintain an Internet connection before a disconnect is forced.”

SO ORDERED this 30th day of August, 2013.


UNITED STATES DISTRICT JUDGE

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

Comcast IP Holdings I, LLC,

Plaintiff,

v.

**Sprint Communications Company L.P.,
Sprint Spectrum L.P., and Nextel
Operations, Inc.,**

Defendants.

Civil Action No. 12-205-RGA

MEMORANDUM ORDER

Presently before the Court is the issue of claim construction of the term “parsing,” found in U.S. Patent Nos. 7,012,916 (“the ‘916 patent”) and 8,204,046 (“the ‘046 patent”). Until recently, the parties had agreed that “parsing” had a plain and ordinary meaning. Now, Comcast proposes that it be construed as, “An automated process of determining the syntactic structure of a language unit by decomposing it into more elementary subunits and establishing the relationships among the subunits.” Sprint proposes that it be construed as, “A process of analyzing a string according to a set of rules of a grammar. Parsing does not include mapping a telephone number to a domain name, nor a lookup from a telephone number to a domain name, and it does not refer to any predetermined association between a number string and a domain name. Parsing is also not the same as translating.” There are two main issues regarding the proposed constructions. The first is whether parsing requires an automated process. The second is whether Sprint’s proposed negative limitations are supported by the prosecution history. The Court heard oral argument on August 20, 2014.

Comcast argues that the plain meaning of “parsing” in the context of computer science requires an automated process. Sprint replies that where the patentee wanted to make clear that parsing was done by a computer, the claims specifically stated that the computer did the parsing. Sprint points to claim 90 of the ‘046 patent and claim 45 of the ‘916 patent to illustrate this difference. Claim 90 covers:

A method, comprising:
 forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name; and
 supplying by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

(‘046 patent at 45:65-46:7). Claim 45 covers:

A method of accessing communications data for contacting a target entity, said method comprising:
 forming, from a number string identifying the target entity, a domain name by a process including parsing at least a substantial portion of the number string into at least a part of said domain name;
 supplying the domain name formed to a DNS-type database system and receiving back a resource record including an URI for locating communications data associated with the domain name; and
 using the URI received back to access said communications data.

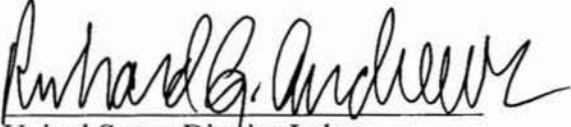
(‘916 patent at 36:55-67). While only claim 90 specifically requires that a computer perform the parsing, it is clear that parsing is an automated process. The reference to a computer in claim 90 is merely to make clear that the same device performs both steps of the method. In fact, Sprint does not contest that the invention relates to computing. The fact that the patentee mentioned a computer in one claim does not mean that the other claims do not require computers. I therefore construe “parsing” to require an automated process.

Turning to Sprint’s negative limitations, I find that they are not clear disavowals and would confuse, rather than help, the jury. As Comcast’s counsel stated at oral argument, “the

presence of parsing should turn on the meaning of parsing.” I agree. Both parties seem to agree that parsing is not a one to one relationship and is not a lookup table.¹ But as Sprint repeatedly pointed out during oral argument, parsing must include “analyzing a string according to a set of rules of a grammar.” Comcast does not disagree with that definition of parsing (as stated at oral argument). I see no reason to add in negative limitations. I do not dispute that the patentee described things that are not by themselves parsing during the prosecution of the patent in order to elucidate the meaning of “parsing.” I do not think, however, that use of those negative limitations rose to the level of a clear disavowal of any claim scope contained within the plain meaning of parsing. I also note that the cited references are not in the record, and Sprint therefore has not shown any disclaimer, assuming that there were a disclaimer.

Additionally, the presence of multiple negative limitations has the distinct possibility of confusing the jury. The fact that the accused system might at some point employ mapping or a lookup table does not mean that it does not parse. I therefore construe “parsing” as, “An automated process of analyzing a string according to a set of rules of a grammar.”

Entered this 22nd day of August, 2014.


United States District Judge

¹ Comcast made the caveat that a lookup table might be created by parsing.

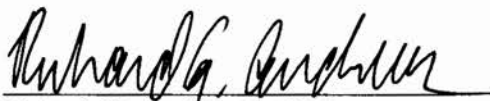
IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

COMCAST IP HOLDINGS I, LLC,	:	
	:	
Plaintiff,	:	
	:	
v.	:	Civil Action No. 12-205-RGA
	:	
SPRINT COMMUNICATIONS	:	
COMPANY, L.P., et al.,	:	
	:	
Defendants.	:	

JUDGMENT

This 22nd day of October 2014, the Court having held a jury trial, and the jury having rendered a verdict, pursuant to Fed. R. Civ. P. 58(b)(2), IT IS HEREBY ORDERED that:

Judgment in the amount of \$7,500,000.00 is entered in favor of Plaintiff Comcast IP Holdings I, LLC and against Defendants Sprint Communications Company L.P., Sprint Spectrum L.P., and Nextel Operations, Inc. on Counts 1, 5, and 7 of the Second Amended Complaint (D.I. 59).


United States District Judge

UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

COMCAST IP HOLDINGS I, LLC,)	
)	
Plaintiff,)	
)	Case No.: 1:12-cv-0205-RGA
v.)	
)	
SPRINT COMMUNICATIONS COMPANY)	
L.P.; SPRINT SPECTRUM L.P.; and)	
NEXTEL OPERATIONS, INC.)	
)	
Defendants.)	
)	
)	

~~[MOTION PROPOSED]~~ JUDGMENT

This action came before the Court for a trial by jury on October 6, 2014. The issues have been tried, and the jury rendered its verdict on October 15, 2014. The verdict was accompanied by a verdict form (D.I. 334). A copy of the public version of that verdict form (D.I. 335) is attached hereto. On October 22, 2014, the Court entered judgment pursuant to Rule 58(b)(2) in favor of Plaintiff Comcast IP Holdings I, LLC (“Comcast”) and against Defendants Sprint Communications Company L.P., Sprint Spectrum L.P., and Nextel Operations, Inc. (collectively, “Sprint”) on Counts 1, 5, and 7 of the Second Amended Complaint (D.I. 342). Before trial, Defendant Sprint filed a Motion for Partial Summary Judgment that the Asserted Claims of U.S. Patent No. 6,873,694 are Invalid Under 35 U.S.C. § 101 (D.I. 151), and the Court granted that Motion on July 16, 2014 (D.I. 292), invalidating Claim 21 of U.S. Patent No. 6,873,694.

Therefore, in accordance with the jury’s verdict and the Court’s orders, IT IS ORDERED AND ADJUDGED:


That judgment be and is hereby entered in favor of Plaintiff Comcast and against Defendant Sprint that Sprint directly infringes Claim 45 of U.S. Patent No. 7,012,916, Claims 90

and 113 of U.S. Patent No. 8,204,046, and Claims 1, 13 and 27 of U.S. Patent No. 8,170,008, as set out in the attached verdict form (D.I. 335);

That damages be and are hereby awarded in favor of Plaintiff Comcast and against Defendant Sprint for the aforementioned infringement in the amount of \$7.5 million, as set out in the attached verdict form (D.I. 335); and

That judgment be and is hereby entered in favor of Defendant Sprint and against Plaintiff Comcast that Claim 21 of U.S. Patent No. 6,873,694 is invalid under 35 U.S.C. § 101 as set out in the Court's order (D.I. 292).

This Judgment is subject to modification following the Court's consideration of the parties' post-trial motions, if any.


UNITED STATES DISTRICT JUDGE

Dated: November 14, 2014

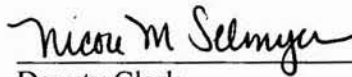

Deputy Clerk

EXHIBIT A

Filed thus 15th Day of October, 2015, In open Court -KJY
IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

COMCAST IP HOLDINGS I, LLC,

Plaintiff,

v.

SPRINT COMMUNICATIONS COMPANY
L.P.; SPRINT SPECTRUM L.P.; and
NEXTEL OPERATIONS, INC.,

Defendants.

C.A. No. 12-205-RGA (CJB)

DEMAND FOR JURY TRIAL

VERDICT FORM

In answering these questions, you are to follow all of the instructions I have given you in the Court's charge. As used herein, "Comcast" means Comcast IP Holdings I, LLC. As used herein, "Sprint" means Sprint Communications Company L.P.; Sprint Spectrum L.P.; and Nextel Operations, Inc.

I. ALLEGED INFRINGEMENT BY SPRINT

Question 1: Did Comcast prove by a preponderance of the evidence that Sprint infringes one or more of Comcast's patent claims? (A "yes" answer is a decision in favor of Comcast, and a "no" answer is a decision in favor of Sprint.)

☒ Yes ☐ No

If you answered "Yes" to Question 1, check the appropriate box(es) below to indicate the claim(s) Sprint infringes and the accused call flows that infringe the respective claim(s). (A checked box is a decision in favor of Comcast and a box left blank is a decision in favor of Sprint.)

'916 Patent

	Claim 45
When an SMI subscriber on a CDMA mobile handset makes a call to another SMI subscriber or a Google Voice subscriber	<input checked="" type="checkbox"/>
When a Google Voice subscriber makes a call to an SMI subscriber	<input checked="" type="checkbox"/>
When a Sprint subscriber using an Airave 2 device makes a call to an SMI subscriber or a Google Voice subscriber	<input checked="" type="checkbox"/>

'046 Patent

	Claim 90	Claim 113
When an SMI subscriber on a CDMA mobile handset makes a call to another SMI subscriber or a Google Voice subscriber	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
When a Google Voice subscriber makes a call to an SMI subscriber	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
When a Sprint subscriber using an Airave 2 device makes a call to an SMI subscriber or a Google Voice subscriber	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

'008 Patent

	Claim 1	Claim 13	Claim 27
When an SMI subscriber on a CDMA mobile handset makes a call to any party except for a Sprint subscriber that is not an SMI or Google Voice user	✓	✓	✓
When a Google Voice subscriber makes a call to any party except for a Sprint subscriber that is not an SMI user	✓	✓	✓
When a Sprint subscriber on Sprint's CDMA network makes a call to a user of an Airave 2 device	✓	✓	✓

II. DAMAGES

Only answer Question 2 if you have found that Sprint has infringed one or more claims.

Question No. 2: What is the amount of damages you have determined for infringement?

\$ 7.5 million

You have now reached the end of the verdict form and should review it to ensure it accurately reflects your unanimous determinations. The Foreperson should then sign and date the verdict form in the spaces below and notify the Court Security Officer that you have reached a verdict. The Foreperson should retain possession of the verdict form and bring it when the jury is brought back into the courtroom.

DATED: 10/15, 2014

Foreperson

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

COMCAST IP HOLDINGS I, LLC,

Plaintiff,

v.

SPRINT COMMUNICATIONS COMPANY
L.P, SPRINT SPECTRUM L.P.; and
NEXTEL OPERATIONS, INC.,

Defendants.

No. 12-cv-0205-RGA

Memorandum Opinion

Matthew Lehr, Esq. (argued), Davis Polk & Wardwell, LLP, Menlo Park, CA; Anthony Fenwick, Esq. (argued), Davis Polk & Wardwell LLP, Menlo Park, CA; Arthur G. Connolly, III, Connolly Gallagher LLP, Wilmington, DE; for the Plaintiff.

Brian C. Riopelle, Esq. (argued), McGuire Woods, Richmond, VA; David E. Finkelson, Esq. (argued), McGuire Woods, Richmond, VA; Rachelle H. Thompson, Esq. (argued), McGuire Woods, Raleigh, NC; Richard Herrmann, Morris James LLP, Wilmington, DE; for the Defendants.

August 10, 2015


ANDREWS, UNITED STATES DISTRICT JUDGE:

Before this Court is Defendants' combined Motion for Judgment as a Matter of Law and Motion for New Trial. (D.I. 352). It has been fully briefed. (D.I. 353, 362, 368). Plaintiff has also moved for post-trial relief (D.I. 355), and that motion is also fully briefed. (D.I. 356, 360, 367). In connection with its request for post-trial relief, Plaintiff has also submitted a letter amending its claimed prejudgment interest. (D.I. 374). The Court has heard oral argument on these post-trial matters. (D.I. 375).

After a 6-day jury trial in October 2014, judgment was entered for Plaintiff against Defendants consistent with the jury's verdict for infringing claim 45 of U.S. Patent No. 7,012,916 ("916 patent"), claims 90 and 113 of U.S. Patent No. 8,204,046 ("046 patent"), and claims 1, 13, and 27 of U.S. Patent No. 8,170,008 ("008 patent"). (D.I. 351). Damages were awarded to Plaintiff for \$7.5 million. *Id.* At the close of Plaintiff's case-in-chief, Defendants moved for judgment as a matter of law on non-infringement of the asserted patents as well as damages. (D.I. 329 at 115-16). Defendants renewed this motion after the close of evidence, which the Court denied. (D.I. 330 at 139-41).

The patents at issue broadly relate to using internet technology to route calls through a network. (D.I. 328 at 230). Both parties are in agreement about what call flows are accused of infringement. (D.I. 353 at p. 7; D.I. 362 at pp. 4-5). Comcast contended that Sprint infringed its patents-in-suit when providing telephony service to subscribers using (1) Sprint Mobile Integration ("SMI"), (2) Google Voice, or (3) Sprint's Airave 2 device. (D.I. 332 at 15). For the '008 patent's asserted claims, Comcast accused three call flows of infringing: (1) when an SMI subscriber on a CDMA mobile handset makes a call to any party except for a Sprint subscriber that is not an SMI or Google Voice user, (2) when a Google Voice subscriber makes a call to any

party except for a Sprint subscriber that is not an SMI user, and (3) when a Sprint subscriber on Sprint's CDMA network makes a call to a user of an Airave 2 device. (D.I. 335 at 3). Comcast also alleged that three call flows infringed the asserted claims of both the '916 and '046 patents: (1) when an SMI subscriber on a CDMA mobile handset makes a call to another SMI subscriber or a Google Voice subscriber, (2) when a Google Voice subscriber makes a call to an SMI subscriber, and (3) when a Sprint subscriber using an Airave 2 device makes a call to an SMI subscriber or a Google Voice subscriber. (D.I. 335 at 2).

For the following reasons, Sprint's Renewed Motion for Judgment as a Matter of Law and Motion for a New Trial (D.I. 352) is **DENIED**, and Comcast's motion for post-trial relief (D.I. 355) is **GRANTED**.

I. LEGAL STANDARD

A. JMOL

Judgment as a matter of law is appropriate if "the court finds that a reasonable jury would not have a legally sufficient evidentiary basis to find for [a] party" on an issue. Fed. R. Civ. P. 50(a)(1). "Entry of judgment as a matter of law is a 'sparingly' invoked remedy, 'granted only if, viewing the evidence in the light most favorable to the nonmovant and giving it the advantage of every fair and reasonable inference, there is insufficient evidence from which a jury reasonably could find liability.'" *Marra v. Phila. Hous. Auth.*, 497 F.3d 286, 300 (3d Cir. 2007) (citation omitted).

"To prevail on a renewed motion for JMOL following a jury trial, a party must show that the jury's findings, presumed or express, are not supported by substantial evidence or, if they were, that the legal conclusion(s) implied [by] the jury's verdict cannot in law be supported by those findings." *Pannu v. Iolab Corp.*, 155 F.3d 1344, 1348 (Fed. Cir. 1998) (alterations in original). "'Substantial' evidence is such relevant evidence from the record taken as a whole as

might be accepted by a reasonable mind as adequate to support the finding under review.”

Perkin-Elmer Corp. v. Computervision Corp., 732 F.2d 888, 893 (Fed. Cir. 1984).

In assessing the sufficiency of the evidence, the Court must give the non-moving party, “as [the] verdict winner, the benefit of all logical inferences that could be drawn from the evidence presented, resolve all conflicts in the evidence in his favor and, in general, view the record in the light most favorable to him.” *Williamson v. Consol. Rail Corp.*, 926 F.2d 1344, 1348 (3d Cir. 1991). The Court may “not determine the credibility of the witnesses [nor] substitute its choice for that of the jury between conflicting elements in the evidence.” *Perkin-Elmer*, 732 F.2d at 893. Rather, the Court must determine whether the evidence supports the jury’s verdict. *See Dawn Equip. Co. v. Ky. Farms Inc.*, 140 F.3d 1009, 1014 (Fed. Cir. 1998); *Gomez v. Allegheny Health Servs. Inc.*, 71 F.3d 1079, 1083 (3d Cir. 1995) (describing standard as “whether there is evidence upon which a reasonable jury could properly have found its verdict”); 9B *Charles Alan Wright & Arthur R. Miller, Federal Practice and Procedure* § 2524 (3d ed. 2008) (“The question is not whether there is literally no evidence supporting the party against whom the motion is directed but whether there is evidence upon which the jury might reasonably find a verdict for that party.”).

B. MOTION FOR A NEW TRIAL

Federal Rule of Civil Procedure 59(a)(1)(A) provides, in pertinent part: “The court may, on motion, grant a new trial on all or some of the issues—and to any party— . . . after a jury trial, for any reason for which a new trial has heretofore been granted in an action at law in federal court” Among the most common reasons for granting a new trial are: (1) the jury’s verdict is against the clear weight of the evidence, and a new trial must be granted to prevent a miscarriage of justice; (2) newly discovered evidence exists that would likely alter the outcome

of the trial; (3) improper conduct by an attorney or the court unfairly influenced the verdict; or (4) the jury's verdict was facially inconsistent. *See Zarow-Smith v. N.J. Transit Rail Operations, Inc.*, 953 F. Supp. 581, 584–85 (D.N.J. 1997).

The decision to grant or deny a new trial is committed to the sound discretion of the district court. *See Allied Chem. Corp. v. Daiiflon, Inc.*, 449 U.S. 33, 36 (1980); *Olefins Trading, Inc. v. Han Yang Chem. Corp.*, 9 F.3d 282, 289 (3d Cir.1993) (reviewing district court's grant or denial of new trial motion under the "abuse of discretion" standard). Although the standard for granting a new trial is less rigorous than the standard for granting judgment as a matter of law—in that the Court need not view the evidence in the light most favorable to the verdict winner—a new trial should only be granted where "a miscarriage of justice would result if the verdict were to stand," the verdict "cries out to be overturned," or where the verdict "shocks [the] conscience." *Williamson*, 926 F.2d at 1352–53.

II. DISCUSSION

A. Judgment as a Matter of Law or New Trial

1. Jury verdict on the '008 patent

At trial, sufficient evidence was offered for a jury to determine that the accused call flows infringed the asserted claims of the '008 patent, and thus, this Court will not grant judgment as a matter of law or a new trial on these grounds. Defendants argue that the jury's verdict rendered three terms in the three asserted claims of the '008 patent superfluous, "call destination," "identifier of a second party," and "second party." (D.I. 353 at p. 12). These terms were not

previously construed by this Court, so the jury was charged with applying their “common meaning.” (D.I. 332 at 18).¹

Claim 1 reads:

A method, comprising: receiving, over a switched telecommunication system, a request; determining, responsive to the request, *a call destination* using domain name system signaling; and initiating a call through the switched telecommunication system between a calling party and *the call destination* that was determined as a result of said domain name system signaling.

(‘008 patent, claim 1) (emphasis added).

Claim 13 reads:

A method, comprising: receiving, over a switched telecommunication system, an indication of a called party; determining, responsive to the indication of the called party, *a call destination* associated with the called party using domain name system signaling; and initiating a call through the switched telecommunication system between a calling party and *the call destination* that was determined as a result of said domain name system signaling.

(‘008 patent, claim 13) (emphasis added).

Claim 27 reads:

A method, comprising: receiving a request from a first party; determining by a computer, responsive to the request, *an identifier of a second party* using domain name system signaling; setting up a call through the switched telecommunication system between the first party and *the second party* that was determined as a result of said domain name system signaling.

(‘008 patent, claim 27) (emphasis added).

There was substantial evidence to support the jury’s determination that Sprint’s accused call flows met the call destination limitation. Sprint argues that Comcast eviscerated the meaning of “call destination” by implying through its expert, Dr. Eric Burger, that a call

¹ The Court construed about ten disputed terms, and the parties reached agreement on a number of other constructions. (D.I. 138). No one requested that the Court construe the terms in the ‘008 patent now at issue. (D.I. 69).

destination is “the place that Sprint routes calls to.” (D.I. 353 at p. 13; *see* D.I. 328 at 357-58 (“Q. And it is your opinion in this case that the call destination that is required by claims 1 and 13 of the ‘008 patent is something other than the receiving phone; is that correct? A. That’s correct. ...Actually, it’s the place that Sprint routes the call to.”)). Thus, Sprint argues it was improper that Comcast’s attorney then summarized this position in closing, arguing that a call destination is “any destination along the way in route of a call” and can be “multiple destinations.” (D.I. 353 at p. 13; D.I. 333 at 46-49). It seems that Sprint wants a call destination to mean the absolute final point in the call flow, or the receiving phone. But a jury could reasonably conclude that there are intermediate “call destinations” before the phone receives the call, or even several “call destinations,” especially relying on the testimony of Dr. Burger at trial. (*See* D.I. 328 at 357-58). Further, Sprint had an opportunity to cross-examine Dr. Burger, and attack his credibility, which Sprint did. (D.I. 328 at 358 (“Q. ...A call destination, in your view, is the final point in a particular service provider’s network to which a call is routed; is that correct? A. Well, I think sometimes I may have misspoke, but it’s where Sprint routes the call.”)). It was the jury’s role to weigh the credibility of Dr. Burger’s testimony and to accept it if persuasive. To accept his testimony does not mean the term was read out of the claims; not every point along the route of a call is a call destination. Beyond relying on Dr. Burger’s testimony, a jury could reasonably look at the claims and conclude that the word “a” before the first use of the term call destination implied that there could be more than one destination, and that the “the” before the second use of the term call destination was the final point in reference to where the calls were routed by Sprint’s networks performing the patented method. (*See* ‘008 patent, claims 1, 13).

Similarly, the jury's verdict did not render the terms "identifier of a second party" and "second party" superfluous in claim 27 of the '008 patent, as Defendants argue. (D.I. 353 at p. 15). Defendants appear to argue essentially that the second party cannot be a third party service provider or an Airave 2 device but must be a dialed party's handset. (D.I. 353 at pp. 15-16). Dr. Burger explained that DNS (domain name system) signaling determines an "identifier of a second party," which could be, for example, the MAC (media access control) address of an Airave II device (D.I. 328 at 252) or an SBC (session border controller) (D.I. 328 at 258), which interfaces between Sprint's network and other networks. (D.I. 328 at 125-26). Sprint cross-examined Dr. Burger about allegedly dropping a requirement that the "identifier of a second party" must "uniquely" identify a second party. (*See, e.g.*, D.I. 328 at 356 ("Q. If the IP address and port of the SBC or [MGC] [media gateway controller] are not in fact unique to any particular carrier or service provider, then you would agree that the SBC and [MGC] do not identify a second party for purposes of claim 27; is that correct? A. No. They still identify the second party. It's how you get to the second party.")). A jury could reasonably rely on Dr. Burger's testimony and find him credible, rejecting the weight of this cross-examination.

Sprint makes broader arguments about how a call cannot exist between a mobile handset and a device such as an Airave 2, SBC and MGC. (D.I. 353 at p. 16). The crux of Defendants' argument is that, because one cannot "talk" to these devices, a "call" as determined by the claims cannot be set up. (D.I. 353 at p. 16). This is a spurious argument. In the context of routing calls through networks, it is unclear why "talking" is a requirement of a "call." Plaintiff rightfully points out that "nothing in the claims requires that a calling party be able to 'talk to' the call destination or the 'second party.'" (D.I. 362 at p. 16). The jury was free to use an ordinary meaning of call that did not require "talking to" the other party. Therefore this argument fails.

Sprint wrongly relies upon Dr. Burger's testimony to demonstrate that the jury verdict is incorrect as a matter of law when considering Dr. Burger's testimony about non-accused call flows. (D.I. 353 at pp. 17-18). As a general proposition, it seems challenging to draw too much from the call flows that Dr. Burger carved out from his infringement examples. (D.I. 328 at 362-63 ("Q. ... You specifically identified them as an exception to your infringing use cases? A. In my summary table, that's correct. Q. You agree they don't infringe? A. I -- I have not given an opinion, so... Q. You have not opined that they infringe; is that correct? A. That's correct.")). Plaintiff says that Dr. Burger agreed that these call flows were exceptions to his infringing call flows, but that only means that they were not call flows he "affirmatively concluded" infringe. (D.I. 362 at p. 17). It seems clear that Dr. Burger did not explicitly determine that they do not infringe. (See D.I. 328 at 362 ("Q. And that's because you agree that a call from an SMI or Google Voice CDMA handset to a Sprint subscriber does not infringe the Comcast patents that are at issue in this case; is that correct? A. Actually, I can't remember whether I did that as a determination. I didn't determine that they did infringe. Q. In fact, you specifically identified them as an exception -- A. As an exception.")).

Sprint wants to read these examples as demonstrating that, because the MGC and SBC are not destinations in these exceptions, they cannot be destinations in the infringing examples. Considering how tentative Dr. Burger was with his non-infringing examples, this is not a reasonable position. In its reply brief, Sprint says that Dr. Burger carved out the unaccused call flows as non-infringing exceptions and that he "specifically explained why they do not infringe," citing a statement from the transcript, "because after routing to the MGC, it continues to another point in Sprint's network for further routing." (D.I. 368 at p. 11). That statement, however, is not Dr. Burger "specifically explain[ing]" the non-infringing exceptions, but a quote from

Sprint's counsel's question, with which Dr. Burger agreed. (D.I. 328 at 363 ("Q. In fact, you specifically singled them out as not infringing, as being exceptions to your infringing use case? A. That's right. Q. You agree that a call from an SMI caller to a regular Sprint subscriber, a Sprint CDMA subscriber, does not infringe the '008 patent because after routing to the MCG, it continues on to another point in Sprint's network for further routing; is that correct? A. That's correct.")). Assuming that this could be characterized as inconsistent with his other testimony, it does not follow that there is not substantial evidence to support the jury's finding that the accused call flows met the call destination limitation. Furthermore, the relevance of Dr. Burger's exceptions is limited because of how tempered Dr. Burger was about why he had not determined that they were infringing.

Finally, Sprint argues that the call destination where a Sprint subscriber on Sprint's CDMA network makes a call to a user of an Airave 2 device does not use domain name system signaling, as required by the claims. (D.I. 353 at pp. 18-19). Sprint essentially argues that the domain name system signaling process only determines the location of a convergence server, not the Airave 2 device. (D.I. 353 at p. 19). But Plaintiff reasonably argues that the MAC address of the Airave 2 device is determined by a process that includes domain name system signaling. (D.I. 362 at p. 18; *see also* D.I. 328 at 115-16). It seems reasonable that a jury could conclude from this evidence that this call flow satisfies the requirement of using domain name system signaling.

2. Jury verdict on the '916 and '046 patents

At trial, sufficient evidence was offered for a jury to determine that the accused call flows infringed the asserted claims of the '916 and '046 patent. Therefore this Court will not grant judgment as a matter of law or a new trial on these grounds. Defendants argue that there is no

legally sufficient evidence to support a jury verdict that satisfied the “parsing” limitation of the claims. (D.I. 353 at pp. 19-22). This Court construed “parsing” as “[a]n automated process of analyzing a string according to a set of rules of a grammar.” (D.I. 298 at 3).

Claim 45 of the ‘916 patent reads:

A method of accessing communications data for contacting a target entity, said method comprising: forming, from a number string identifying the target entity, a domain name by a process including *parsing* at least a substantial portion of the number string into at least a part of said domain name; supplying the domain name formed to a DNS-type database system and receiving back a resource record including an URI for locating communications data associated with the domain name; and using the URI received back to access said communications data.

(‘916 patent, claim 45) (emphasis added).

Claim 90 of the ‘046 patent reads:

A method, comprising: forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including *parsing* at least a portion of the number string into at least a part of said domain name; and supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

(‘046 patent, claim 90) (emphasis added).

Claim 113 of the ‘046 patent reads:

A method, comprising: forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including *parsing* at least a portion of the number string into at least a part of said domain name; supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity; and sending a message to the target entity identified using the URI.

(‘046 patent, claim 113) (emphasis added).

As the Court has said previously, there is sufficient evidence to find that the “parsing” limitation has been met, particularly when considering Dr. Burger’s testimony. (See D.I. 329 at 121 (“So, on the parsing, I’m going to deny the JMOL, and you know, I think Dr. Burger

touched the basis and I think there's enough analysis there, so that [if] the jury finds for him, for Comcast, that will be supported by substantial evidence.")). Dr. Burger provided extensive detailed testimony about how "parsing" occurred:

...So the judge in a Claim Construction Order said that parsing is an automated process of analyzing a string according to a set of rules of a grammar. And in this situation, you see Sprint and, in fact, Mr. Nehme talked about it, and Mr. Riopelle talked about it yesterday about how Sprint analyzes a string.

And you have to understand in computer science, you know, when you think of analyzing, you're thinking of, you know, what people do, but we're talking about computers here. In computer science terms, a string is just any set of characters.

It could be text. But in this case, it happens to be numbers. And it's a set of rules of a grammar. A grammar in computer science is not [like] English with nouns, verbs and subjects. A grammar is just when you have a set of tokens or just a bit of strings, what's the relationship between the tokens?

In Sprint's case, and as Mr. Nehme explained, how they do this step of parsing? They take all of the digits. That's the analyzing part.

And in the relationship between them is put them in the opposite order, put in the dots and then the E164.ARPV.

(D.I. 328 at 261-62).

Defendants had an opportunity to cross-examine Dr. Burger on this testimony, and they argue Dr. Burger gave the opposite testimony on the very same steps. (D.I. 353 at p. 20). But, in the testimony cited, Defendants tried to discredit Dr. Burger by pinning him down on which step or steps he was asserting were parsing. (*See, e.g.*, D.I. 328 at 340 ("Q. And it's your testimony, correct, Dr. Burger, that the first step is parsing; correct? A. The validating the number, that's – this is, you know, where it gets interesting because a number of these steps taken alone would be parsing. And putting them altogether is definitely parsing. So it's hard to say: Is this step a parsing step or that because it's all part of that analysis.")). Even if Sprint's expert, Dr. M. Ray Mercer, testified that parsing does not occur, the jury was free to weigh his testimony against Dr. Burger's and the rest of the evidence. (*See, e.g.*, D.I. 329 at 242 ("Q. ... So your testimony is that example right there is an example of what ENUM does; correct? A. Yes. Q. And your

testimony is that is not parsing? A. That's exactly right.")). There is sufficient evidence for the jury to have found that parsing occurred, and therefore the jury verdict cannot be disturbed.

3. The damages award

Defendants argue that the Court should vacate the damages award because Comcast's expert Carla Mulhern's methodology lacked sufficient analysis and evidentiary support. (D.I. 353 at pp. 22-24). Ms. Mulhern testified that Comcast is entitled to damages of \$16.5 million for Sprint's infringement of all three patents. (D.I. 329 at 31). More specifically, she opined that Comcast was entitled to damages of \$15 million for the '008 patent, \$1 million for the '916 patent, and \$500,000 for the '046 patent. (D.I. 329 at 31). The jury awarded damages of \$7.5 million for all three patents. (D.I. 351).

Ms. Mulhern's analysis was not a "black box" as Defendants argue, nor did she provide "no explanation" for how she reached her figures. (D.I. 353 at pp. 22-23). Defendants argue that her figures were not based on any analysis except for certain *Georgia-Pacific* factors affecting her figures upwards or downwards. (D.I. 353 at pp. 22-23). Defendants argue that Ms. Mulhern's analysis was based only on her experience and was deficient because it did not use a mathematical formula. (D.I. 353 at p. 24). Defendants rely on *GPNE Corp. v. Apple, Inc.*, 2014 WL 1494247, at *5 (N.D. Cal. April 16, 2014), but in that case there was no methodology beyond the expert's "30 years of experience," no economic analysis, nor sufficient facts or data. Here, however, that is not the case. Ms. Mulhern started her analysis with a 2008 patent transaction that Comcast valued "north of \$4 million," but for which Comcast ultimately paid \$2.5 million. (D.I. 329 at 41-42). She then compared that agreement to a hypothetical negotiation, explaining that factors related to uncertainty at the time would lead to a substantial discount on the purchase price of the transaction. (D.I. 329 at 43-44). She opined that a

multiplier of more than three would be appropriate in this case. (D.I. 329 at 45). Then, using the *Georgia-Pacific* factors, Ms. Mulhern adjusted her figure to arrive at her opined damage award. (See, e.g., D.I. 329 at 53-66).

Defendants are correct that Ms. Mulhern did not offer a mathematical formula for her result. (See D.I. 329 at 79 (“Q. You don’t have any numerical calculations or arithmetic? You haven’t done any analysis that gets you to, for example, the \$16.5 million on the ‘008, no numerical calculation that gets you to the ‘008 damages, do you? A. I mean, I think if what you’re asking is -- I can agree with you that there’s no mathematical formula that gets me to that -- spits out the \$16.5 million number. That’s right.”)). That does not mean that her opinion was not based on a sound methodology. The hypothetical negotiation to determine a reasonable royalty can involve some approximation. See *LaserDynamics, Inc. v. Quanta Computer, Inc.*, 694 F.3d 51, 76 (Fed. Cir. 2012) (“The hypothetical negotiation ... necessarily involves an element of approximation and uncertainty.”) (internal quotation marks omitted) (citations omitted). It was Defendants’ option to raise the lack of a formula on cross-examination, as they did, and the jury was free to discount Ms. Mulhern’s analysis, as it appears the jury did. The jury awarded \$7.5 million, not \$16.5 million. The lack of a mathematical formula, when there is other analysis, cannot, alone, be grounds for excluding Ms. Mulhern’s methodology.

The jury was not limited to Ms. Mulhern’s analysis and could have easily arrived at its \$7.5 million verdict relying on Sprint’s expert Dr. Debra Aron’s testimony² or other evidence, something this Court noted at trial. (D.I. 330 at 140-41 (“... if the jury were to determine that Ms. Mulhern’s analysis was not a sufficient basis to prove her theories by a preponderance of the

² Credibility is not for the Court. Nevertheless, I have now seen Dr. Aron testify in two trials, both times with a high degree of credibility.

evidence, I think the jury could come up with its own calculation of a reasonable royalty, and, in fact, I believe Dr. Aron pretty much invited them to do so, and so I don't think that's grounds for a JMOL.")). It is strange that Defendants argue that there is no evidence for a damages verdict when their own expert explicitly invited the jury to rely on her testimony alone to determine damages. (D.I. 330 at 71 ("I think that I have provided the jury with an anchor upon which to consider a reasonable royalty. I have provided them, I hope, a way of understanding what a reasonable royalty is, and I've provided them, I think, a perspective on the evidence in the case that I would hope they would keep in mind and consider as to what a reasonable royalty would be if they were to decide that they should come up with one.")). Indeed, Dr. Aron provided extensive analysis of how the 2008 patent transaction of \$2.5 million could be used to determine an appropriate damages award. (*See, e.g.*, D.I. 330 at 47-49). In light of this testimony, it seems more than reasonable that a jury could, and probably did, rely on Dr. Aron's testimony alone to arrive at the \$7.5 million verdict.

Sprint's arguments about excluding Comcast's expert Ms. Mulhern also fail. As the Court explained at the hearing for these motions, it is my opinion that the time to bring a *Daubert*-type theory to exclude what is disclosed in the expert's reports and deposition is before trial. (D.I. 375 at 53-54 ("You have a report that's written by Ms. Mulhern, it's got at least two theories in it, you move to exclude one, I grant it, then we go to trial, she testifies as to the [other] one, you cross-examine her on it, you have Dr. Aron point out some weaknesses. And now, in essence, you're saying, oh, by the way, Judge, on *Daubert* she should be out. It's kind of late.")).

Nor do I think Ms. Mulhern violated this Court's order preventing her from testifying on the "income approach" in violation of the entire market value rule. (*See* D.I. 324 at 1-2). "The

entire market value rule allows for the recovery of damages based on the value of an entire apparatus containing several features, when the feature patented constitutes the basis for customer demand.” *Lucent Technologies, Inc. v. Gateway, Inc.*, 580 F.3d 1301, 1336 (Fed. Cir. 2009) (citations omitted) (internal quotation marks omitted). Sprint tries to demonstrate that Ms. Mulhern opined that the patented technology drove customer demand, but she said no such thing. (See D.I. 353 at pp. 26-27; *see, e.g.*, D.I. 329 at 81 (“Q. You agree, right, that feature has nothing to do with the patented technology at issue in this case; right? A. ... the patents are necessary to complete calls in the SMI service generally, which is necessary for Sprint to be able to offer the service and that feature.”)). Stating that patents are necessary to offer a service is not the same thing as making an income approach argument, where recovery would be based on the value of the entire service. Nor is it the same as saying that a specific feature drives customer demand. Ms. Mulhern also did not rely on revenues or lost profits. Therefore, Ms. Mulhern’s testimony was not improper.

4. A new trial

Sprint moves for a new trial incorporating the same arguments it has already made about the weight of the evidence. (D.I. 353 at p. 27). For the same reasons mentioned above, this argument fails.

Sprint also argues that Comcast’s experts, Dr. Burger and Ms. Mulhern, exceeded the scope of their expert reports. This argument fails. For Dr. Burger, Sprint argues that he identified one call flow as “when a Sprint subscriber on Sprint’s CDMA network makes a call to a user of an Airave 2 device” but in his report this call flow ended at a “user” of an Airave 2 device, rather than the device itself. (D.I. 353 at p. 29; *see* D.I. 354 at 9). At trial, Dr. Burger did not say that the call flow ended at the device, even if he determined that was the call destination.

Dr. Burger stated that the call flow did not stop at the Airave 2 device but continued to the subscriber's device. (D.I. 328 at 243-244 (“Q. And why do you say that the Airave II device is a call destination? A. Because it’s where the Sprint network routes the call to. Q. And after the call is routed by the Sprint network to the Airave II device, what happens next to that call? A. The subscriber, hopefully, the subscriber, you know, answers the call. It rings his phone and they get the call.”)). This testimony did not exceed the scope of Dr. Burger’s expert report.

As for Ms. Mulhern’s testimony, Sprint argues that she testified for the first time about aspects of Factor 6 of the *Georgia-Pacific* analysis, related to ancillary financial benefits of the accused products. (D.I. 353 at p. 30). According to Sprint, Ms. Mulhern opined on the “upward impact” on the royalty for these products because of the ancillary benefits. (D.I. 353 at p. 30; *see, e.g.*, D.I. 329 at 66 (“Factors 6 and 8, when we considered the benefit to Sprint of selling the accused – of the ability to [sell] accused products, suggests an upward impact on the royalty because it suggests that Sprint benefited in the ways that we’ve just been discussing through the ability to sell these accused products.”)). Sprint says that Ms. Mulhern’s report contains no explanation of these ancillary benefits absent the Income Approach, which this Court struck. (D.I. 324 at 2; D.I. 354 at 25-26). The Court agrees with Comcast that Ms. Mulhern did address the ancillary benefits in her report, separate from her Income Approach analysis. (D.I. 354 at 26 (“In this case, the provision of the accused products does, in fact, drive ancillary revenues and profits associated with other products and services.”)). Therefore Ms. Mulhern’s testimony did not exceed the scope of her expert report.

B. POST-TRIAL RELIEF

Comcast has separately moved for post-trial relief for (1) prejudgment interest, (2) postjudgment interest on the jury's \$7.5 million damages award, and (3) an ongoing royalty for continued use of the patented technology. (D.I. 355 at 1). Initially Comcast sought prejudgment interest of \$2,921,451, but Comcast later submitted a letter adjusting its request for prejudgment interest to \$1,691,640, which corrected the earlier request to adjust the date that interest would begin accruing after the hypothetical negotiation. (D.I. 374). At the hearing, Sprint's counsel said there was no objection with the math but purely the underlying methodology. (D.I. 375 at 72). At the hearing, Sprint's counsel also indicated that post-judgment interest was not in dispute. (D.I. 375 at 71).

The parties basically disagree about whether prejudgment interest should be assessed from the time of the hypothetical negotiation in late 2006 (D.I. 356 at 10) or whether it should be calculated from the time of first infringement of each patent until judgment. (D.I. 360 at 14-15).

"As a rule, prejudgment interest should be awarded under 35 U.S.C. § 284 absent some justification for withholding such an award." *Whitserve, LLC v. Computer Packages, Inc.*, 694 F.3d 10, 36 (Fed. Cir. 2012) (citations omitted) (internal quotation marks omitted). "The rate of prejudgment interest and whether it should be compounded or un compounded are matters left largely to the discretion of the district court. In exercising that discretion, however, the district court must be guided by the purpose of prejudgment interest, which is to ensure that the patent owner is placed in as good a position as he would have been had the infringer entered into a reasonable royalty agreement." *Bio-Rad Labs., Inc. v. Nicolet Instrument Corp.*, 807 F.2d 964, 969 (Fed. Cir. 1986) (citations omitted) (internal quotation marks omitted).

The Court agrees with Comcast that the measure of damages must be from the time of the hypothetical negotiation to trial for all the asserted patents. (*See* D.I. 356 at 12). Sprint points

out that two of the patents did not issue until several years after the hypothetical negotiation—the ‘916 patent issued in 2006, but the ‘046 and ‘008 patents issued only in 2012. (D.I. 360 at 8). Even so, both sides’ experts opined that the reasonable royalty should be a lump sum resulting from a single hypothetical negotiation that took place in 2006. (*See, e.g.*, D.I. 330 at 82 (“Q. You agree with Ms. Mulhern that if the ‘916 patent or any combination of patents that includes the ‘916 is found to be infringed, then the appropriate date for the hypothetical negotiation for all infringed patents in this case is late 2006; correct? A. Yes, that’s right.”)). A hypothetical negotiation in 2006 would include each of the asserted patents, even if they issued later, and the resulting negotiation would include the fact that two of the patents issued later. *See ResQNet.com, Inc. v. Lansa, Inc.*, 594 F.3d 860, 872 (Fed. Cir. 2010) (for reasonable royalty calculations, a “district court may also consider the panoply of events and facts that occurred thereafter and that could not have been known to or predicted by the hypothesized negotiators.”) (citations omitted) (internal quotation marks omitted). Further, as Comcast argues, this Court cannot reasonably infer from the jury’s verdict what portion of the damages is attributed to each patent especially as the jury did not adopt Ms. Mulhern’s proposed damage award. (D.I. 367 at pp. 3-4).

Sprint points out that Ms. Mulhern identified different first infringement dates for each of the three patents, but she did so only to explain how the hypothetical negotiation date would be pushed to a later date if the jury found that some, but not all three patents, were infringed. (D.I. 329 at 35-36 (“If Sprint is not found to have infringed the ‘916 patent, but only either the ‘046 or ‘008 patents, which issued much later, then the hypothetical negotiation would occur in 2012 at the time of first infringement of those patents.”)). But the jury found that Sprint infringed all three patents, so the hypothetical negotiation date would be 2006. Sprint’s new theory to

segment the damages by determining when each patent issued (and therefore when each patent was first infringed) is novel but not grounded in any evidence offered at trial.³

The Court agrees with Comcast that the appropriate interest rate is the prime rate, compounded quarterly. *See XpertUniverse, Inc. v. Cisco Sys., Inc.*, 2013 WL 6118447, at *11 (D. Del. Nov. 20, 2013) (“The Court has broad discretion to select the pre-judgment interest rate to be applied, and the Federal Circuit has held that application of the prime rate is appropriate even if there is no evidence that the patent holder borrowed at the prime rate.”) (citations omitted).

Therefore, this Court will award Comcast the prejudgment interest of \$1,691,640 it has requested.⁴

The parties indicated at the hearing that post-judgment interest was not in dispute (D.I. 375 at 71), and because this Court will grant Comcast’s request for prejudgment interest, it will also grant post-judgment interest.

This Court also grants Comcast an ongoing royalty. “In most cases, where the district court determines that a permanent injunction is not warranted, the district court may wish to

³ Would a hypothetical negotiation in 2006 for a patent not issued until 2012 result in an agreement to make a lump sum payment in 2006 or in 2012? If the payment were going to be in 2006, the licensor ought to get prejudgment interest. I would imagine that if two sophisticated parties made an arms-length agreement that one would pay the other a lump sum for some event that would certainly happen far into the future, they would negotiate some built-in adjustment to account for inflation. They are agreeing to pay the fair market value, and both parties would want the future payment to reflect the fair-market value at the time. Under either scenario, as applied to this case, an award of prejudgment interest is appropriate.

⁴ The Court does not find that Comcast’s supplemental letter changes the analysis. (D.I. 374). That is, even though Comcast adjusted the date interest would accrue from the hypothetical negotiation, Comcast still grounded its prejudgment interest analysis on the single hypothetical negotiation, which occurred in 2006, not multiple dates including 2012 when the ‘046 and ‘008 patents issued.

allow the parties to negotiate a license amongst themselves regarding future use of a patented invention before imposing an ongoing royalty. Should the parties fail to come to an agreement, the district court could step in to assess a reasonable royalty in light of the ongoing infringement.” *Paice LLC v. Toyota Motor Corp.*, 504 F.3d 1293, 1315 (Fed. Cir. 2007). Sprint argues that Ms. Mulhern testified that the lump sum license she relied on was to provide compensation for all infringement, and that she did not offer a time limit for her lump sum. (D.I. 360 at 17). Sprint also raises arguments about the deficiency of Ms. Mulhern’s testimony, some similar to arguments mentioned previously. (D.I. 360 at 20). The lump sum that the jury awarded Comcast only extended through the date of trial, something the experts repeatedly stated. (*See, e.g.*, D.I. 330 at 82 (“Q. And the analysis that both you and Ms. Mulhern have done, you understand is directed to a lump sum royalty through the time of trial; correct? A. That’s what it should be, yes.”)). I do not understand why Sprint argues that the lump sum covered the past, present, and future, when its own expert repeatedly testified that the lump sum only extended to the date of the trial. Comcast also points out that both experts agreed during discovery that the lump sum was only calculated through trial. (D.I. 356 at 8 n.1; *see also* D.I. 356-2 at 3 (“Q. Okay. So would it be fair to say that you agree with Ms. Mulhern that the appropriate form of damages in this case is a lump sum reasonable royalty calculated through trial? A. Yes. Q. And that applies to all the patents? A. Yes.”)). Because the lump sum only extended through trial, Comcast is entitled to an ongoing royalty for infringement.⁵ The Court agrees that Comcast is entitled to an ongoing royalty, and will grant its request for mediation: (1) directing the parties to conduct a mediation to reach agreement on this royalty, and (2) asking the

⁵ I say this with some hesitancy, because I do not recall either expert having any analysis that discounted the lump sum for its coverage of a time period less than the full terms of the patents.

parties to submit the dispute to this Court if mediation fails. I think any mediation could be postponed until resolution of any appeal of this case.

III. CONCLUSION

For the foregoing reasons, Sprint's Renewed Motion for Judgment as a Matter of Law and Motion for A New Trial (D.I. 352) is **DENIED**, and Comcast's motion for post-trial relief (D.I. 355) is **GRANTED**. An appropriate order will follow.

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

COMCAST IP HOLDINGS I, LLC,

Plaintiff,

v.

SPRINT COMMUNICATIONS COMPANY
L.P, SPRINT SPECTRUM L.P.; and
NEXTEL OPERATIONS, INC.,

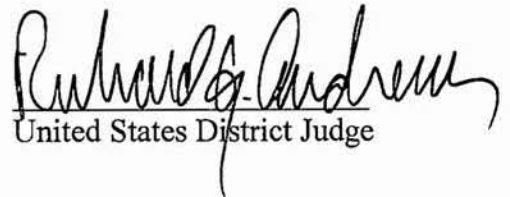
Defendants.

No. 12-cv-0205-RGA


ORDER

For the reasons set forth in the Court's accompanying Memorandum Opinion, Sprint's Renewed Motion for Judgment as a Matter of Law and Motion for A New Trial (D.I. 352) is **DENIED**, and Comcast's motion for post-trial relief (D.I. 355) is **GRANTED**.

IT IS SO ORDERED this 10 day of August 2015.


United States District Judge

- 62 Claims, 14 Drawing Sheets**



US 7,012,916 B2

Page 2

U.S. PATENT DOCUMENTS

5,351,276	A	9/1994	Doll, Jr. et al.	
5,423,003	A	6/1995	Berteau	
5,438,568	A	8/1995	Weisser, Jr.	
5,452,350	A	9/1995	Reynolds et al.	
5,546,452	A	8/1996	Andrews et al.	
5,703,940	A	12/1997	Sattar et al.	
5,708,780	A	1/1998	Levergood et al.	
5,742,905	A	4/1998	Pepe et al.	
5,751,961	A	5/1998	Smyk	
5,799,063	A	8/1998	Krane	
5,799,317	A	8/1998	He et al.	
5,802,146	A	9/1998	Dulman	
5,812,656	A	9/1998	Garland et al.	
5,812,776	A *	9/1998	Gifford	709/217
5,838,682	A	11/1998	Dekelbaum et al.	
5,838,768	A	11/1998	Sumar et al.	
5,870,454	A	2/1999	Dahlen	
5,873,077	A	2/1999	Kanoh et al.	
5,917,817	A	6/1999	Dunn et al.	
5,953,392	A	9/1999	Rhie et al.	
5,958,016	A	9/1999	Chang et al.	
5,966,427	A	10/1999	Shaffer et al.	
6,014,379	A	1/2000	White et al.	
6,014,437	A	1/2000	Acker et al.	
6,021,126	A	2/2000	White et al.	
6,023,724	A	2/2000	Bhatia et al.	
6,026,441	A	2/2000	Ronen	
6,029,203	A	2/2000	Bhatia et al.	
6,067,516	A	5/2000	Levay et al.	
6,084,956	A	7/2000	Turner et al.	
6,097,804	A	8/2000	Gilbert et al.	
6,115,737	A	9/2000	Ely et al.	
6,125,113	A	9/2000	Farris et al.	
6,131,095	A	10/2000	Low et al.	707/10
6,141,413	A	10/2000	Waldner et al.	
6,466,570	B1	10/2002	Low et al.	370/352
6,594,254	B1 *	7/2003	Kelly	370/352
2002/0114324	A1	8/2002	Low et al.	370/352

FOREIGN PATENT DOCUMENTS

EP	0 740 445	10/1996
WO	93/25035	12/1993
WO	94/23523	10/1994
WO	96/09714	3/1996
WO	96/20553	7/1996
WO	96/38018	11/1996
WO	97/16007	5/1997
WO	97/20424	6/1997
WO	97/26749	7/1997
WO	97/32427	9/1997

OTHER PUBLICATIONS

Sevick, et al., "Customers in Driver's Seat: Private Intelligent Network Control Point," *ISS '95 World Telecommunications Congress*, vol. 2, pp. 41-44 (1995).

Yang, et al., "The Design and Implementation of a Service Logic Execution Environment Platform," *Globecom '93*, vol. 3 (Nov. 1993).

Esaki, et al., "Abstraction and Control of Transport Network Resources for Intelligent Networks," *Electronics & Communications in Japan: Part 1—Communications*, vol. 76, No. 1 (Jan. 1993).

Chang, et al., "Rapid deployment of CPE-based telecommunications services," *IEEE Globecom*, vol. 2 pp. 876-880 (Nov. 1994).

Swale, R.P. and D.R. Chesterton, "Distributed intelligence and data in public and private networks," *BT Technology Journal*, vol. 13, No. 2 pp. 95-104 (Apr. 1995).

Comer, et al., "Uniform Access to Internet Directory Services," *Computer Communication Review*, vol. 20, No. 4 pp. 50-59 (Sep. 1990).

Comer, D.E., *Internetworking With TCP/IP vol. 1: Principles, Protocols, and Architecture*, Prentice-Hall, Inc., pp. 329-330 (1991).

Mach2 DNS (Domain Name Services) Module Au-A-11045/97 (Jun. 13, 1995).

Muffett, A., *Proper Care and Feeding of Firewalls*, Sun Microsystems, pp. 1-12 (Nov. 14, 1994).

Berners-Lee, et al., "Hypertext Transfer Protocol—HTTP/1.0," *HTTP Working Group Internet Draft* (Oct. 1995).

Eckardt, T., et al., "On the Personal Communications Impacts on Multimedia Teleservices," *Technical University of Berlin*, pp. 435-449 (1994).

Mockapetris, P., "Domain Names—Concepts and Facilities," *RFC 1034* (Nov. 1987).

Mockapetris, P., "Domain Names—Implementation and Specification," *RFC 1035*, 45 pages (Nov. 1987).

Malamud, C. and Rose, M., "Principles of Operation for the TPC.INT Subdomain: Remote Printing—Technical Procedures," *RFC 1528*, 10 pages (Oct. 1993).

Malamud, C. and Rose, M., "Principles of Operation for the TPC.INT Subdomain: Remote Printing—Administrative Policies," *RFC 1529*, 5 pages (Oct. 1993).

Malamud, C. and Rose, M., "Principles of Operation for the TPC.INT Subdomain: General Principles and Policy," *RFC 1530*, 6 pages (Oct. 1993).

Berners-Lee, T., "Universal Resource Identifiers in WWW," *RFC 1630*, 23 pages (Jun. 1994).

Manning, B. and Colella, R., "DNS NSAP Resource Records," *RFC 1706*, 9 pages (Oct. 1994).

Swale, R., "Virtual Networks of the Future—Converging Public and Private IN," *BT Technology Journal*, pp. 6/1-6/5 (1993).

Swale, et al., "Convergence of Public and Private IN," *BT Technology Journal*, pp. 216-221 (1992).

"Web-On-Call Voice Browser Redefines Access to the Web," *The Netphonic Press Release*, (Mar. 1996).

Atkins, et al., "Integrated Web and Telephone Service Creation," *Bell Labs Technical Journal*, pp. 19-35 (1997).

Tao, et al., "Internet Access via Baseband and Broadband ISDN Gateways," *IEEE International Conference on Computers and Communications*, pp. 485-490 (Apr. 1994).

"Workstation Communications System," *IBM Technical Disclosure Bulletin*, vol. 37, No. 39, (Sep. 1994).

Biala, J., "Mobile Radio and Intelligent Networks," *Principles and Realization of Mobile Communication*, 2nd, revised edition, pp. 42-44, 49-51, 53, 82-83 and 341 (1995).

Everhart, C., et al., "RFC 1183—New DNS RR Definitions," Internet: <<http://www.faqs.org/rfcs/rfc1183.html>> 0 pp. 1-9 (Oct. 1990).

Hvasshovd, Svien-Olaf, et al., "The ClustRa Telecom Database: High Availability, High Throughput, and Real-Time Response," *Proceedings of the 21st VLDB Conference*, pp. 469-477 (1995).

Mockapetris, P. "RFC 1101—DNS Encoding of Network Names and Other Types," Internet: <<http://www.faqs.org/rfcs/rfc1101.html>> pp. 1-12 (Apr. 1989).

US 7,012,916 B2

Page 3

Crocker, D.H., "RFC 822--Standard For the Format of ARPA Internet Text Messages," *RFC 822*Internet: <<http://www.fans.org/rfcs/rfcB22.html>>.

Girod, L., et al., "Requirements for URN Resolution Systems," Internet: <<http://www.potaroo.net/ietf/idref/draft-girod-urn-res-require/>>.

Hamilton, M., "Uniform Resource Identifiers & The Simple Discovery Protocol," *Department of Computer Studies*, Internet: <<http://martinh.net/uris/uris.html>>32 pages total (Jun. 20, 1995).

Hardcastle-Eille, S. E., "RFC 1279 X.500 and Domains," *RFC 1279*Internet: <<http://www.faqs.org/rfcs/rfc1279.html>>pp. 1-12 (Nov. 1991).

Nealling, M., et al., "Uniform Resource Names, ISO OIDs and DNS," Internet: <<http://www.gbiv.com/protocols/uri/drafts/draft-mealling-oid-dns-00.txt>>pp. 1-14 (Nov. 1993).

Mitra, et al., "Uniform Resource Names," Internet: <<http://www.mitra.biz/uri/graft-ictt-uri-resource-names-03.txt>>pp. 1-9 (Nov. 1994).

Rose, M., et al., "RFC 1486--An Experiment In Remote Printing" *RFC 1486*Internet: <<http://www.faqs.org/rfcs/rfc1486.html>>pp. 1-12 (Jul. 1993).

Shafer, K.E. et al., "Urn Services," Internet: <<http://staff.oclc.org/research/publications/shafer/urn/draft-shafer-uri-urn-resolution-00.html>>pp. 1-10 (Jun. 1995).

Shafer, K.E., et al., "Urn Services," Internet: <<http://staff.oclc.org.oclc/research/publications/shafer/urn/draft-ietf-uri-urn-resolution-01.html>>pp. 1-14 (Jul. 1995).

Sollins, K. et al., "RPC 1737--Functional Requirement for Uniform Resource Names," *RFC 1737*, Internet: <<http://www.faqs.org/rfcs/rfc1737.html>>pp. 1-6 (Dec. 1994).

The Urn Implementors, "Uniform Resource Names," *D-Lib Magazine*, Internet: <<http://www.dlib.org/dlib/February96/02arms.html>>pp. 1-6 (Feb. 1996).

* cited by examiner




FIG. 1
(PRIOR ART)




FIG. 2
(PRIOR ART)

U.S. Patent

Mar. 14, 2006

Sheet 2 of 14

US 7,012,916 B2

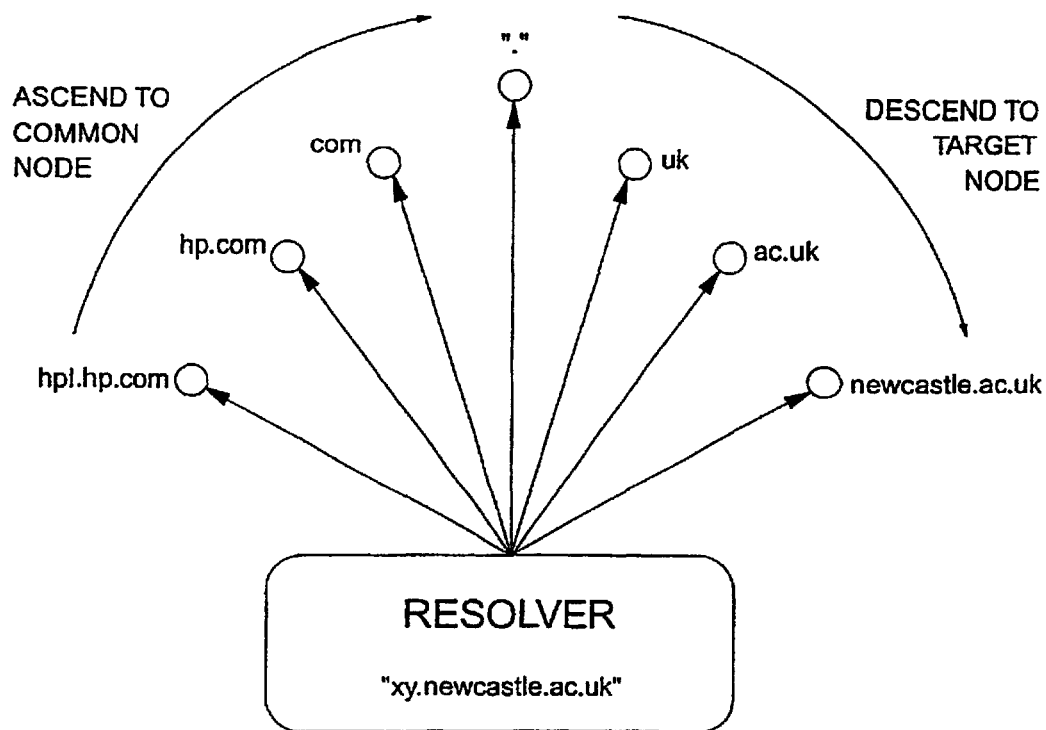


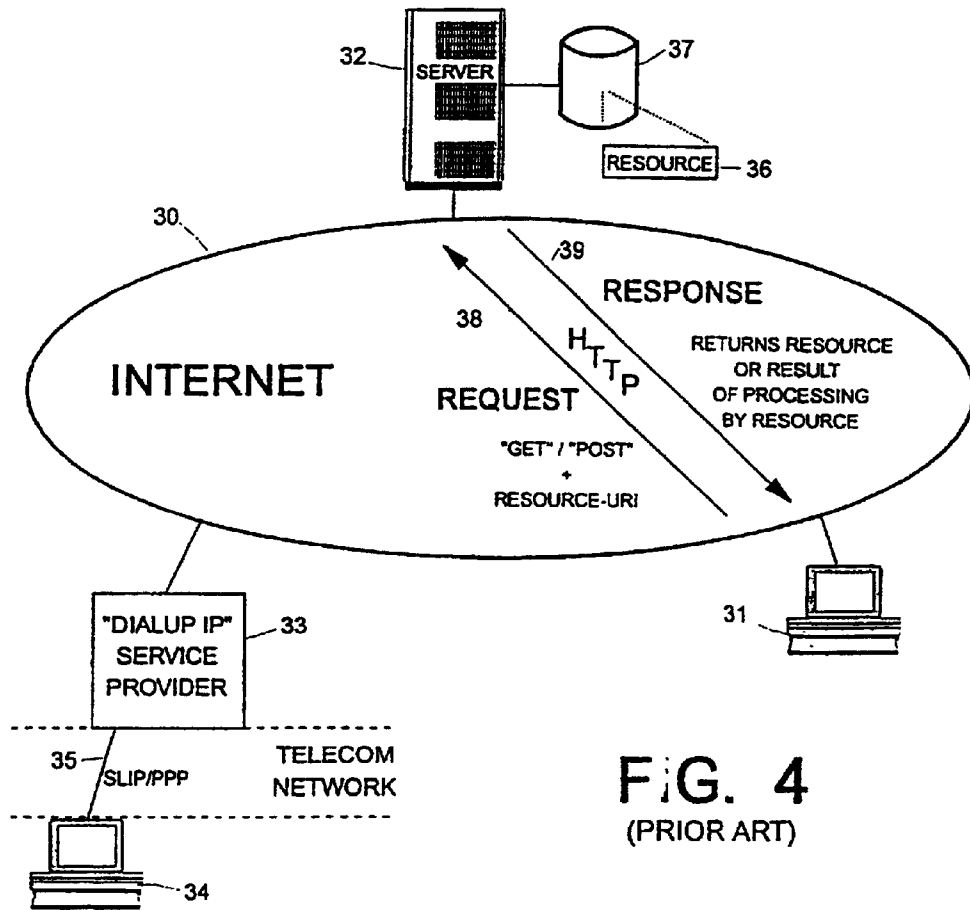
FIG. 3
(PRIOR ART)

U.S. Patent

Mar. 14, 2006

Sheet 3 of 14

US 7,012,916 B2



http : //www.hp.com / Products.html

SCHEME HOST LOCATION ABSOLUTE PATH

FIG. 5
(PRIOR ART)

U.S. Patent

Mar. 14, 2006

Sheet 4 of 14

US 7,012,916 B2




FIG. 6

A77

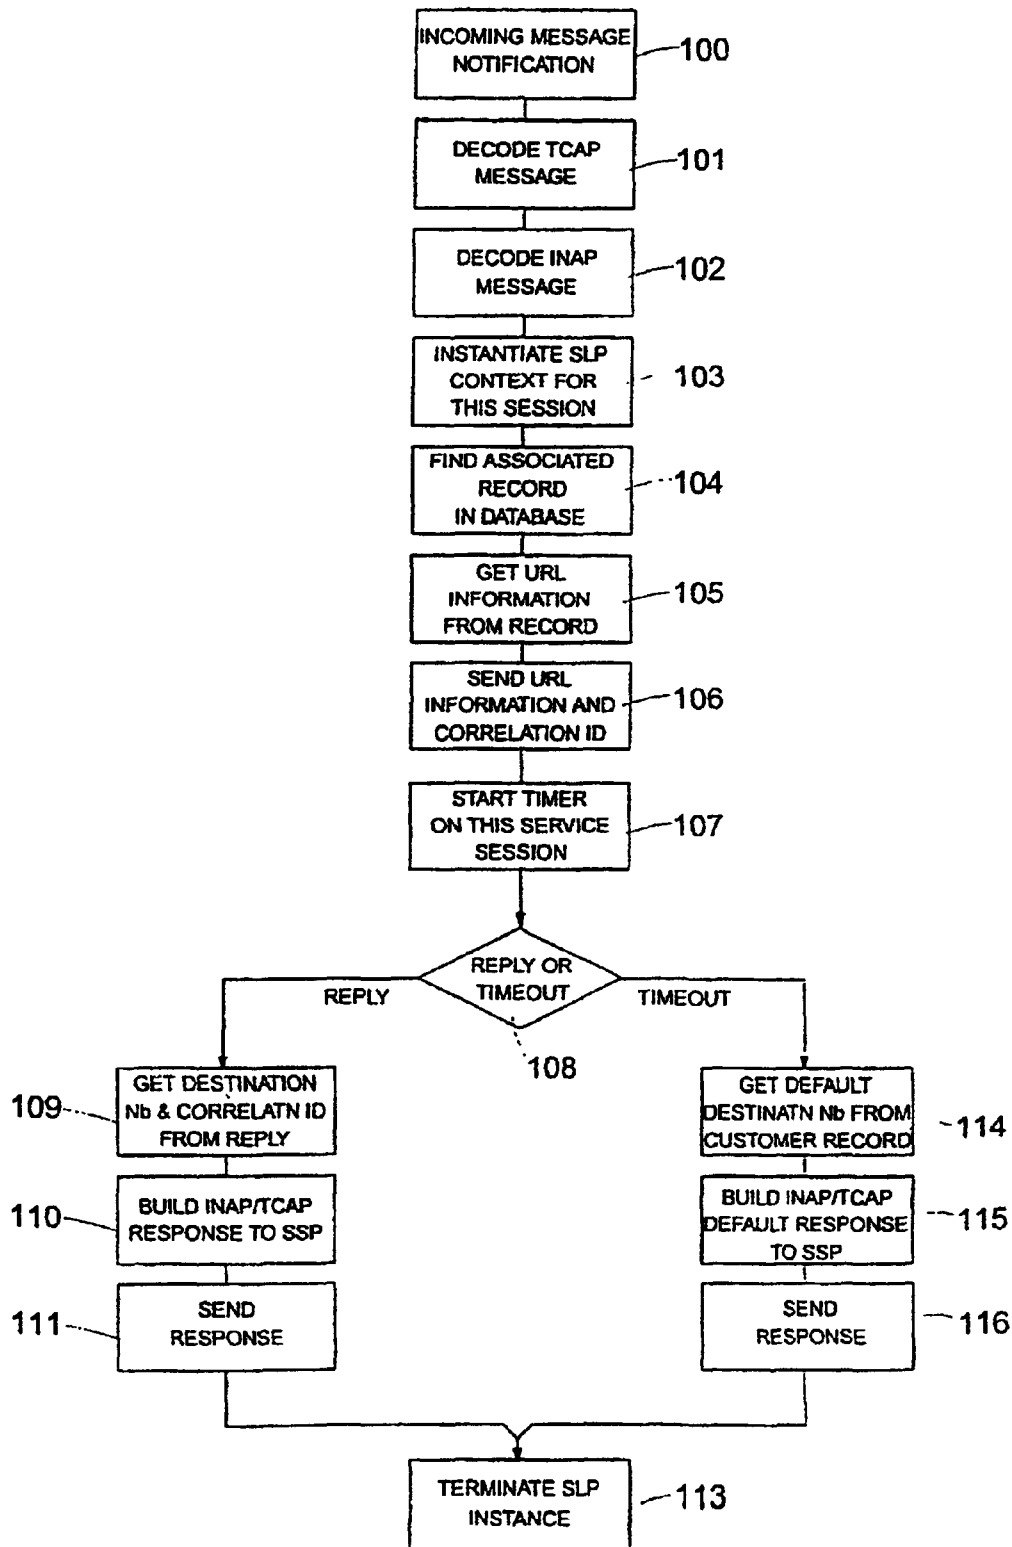


FIG. 7




FIG. 8

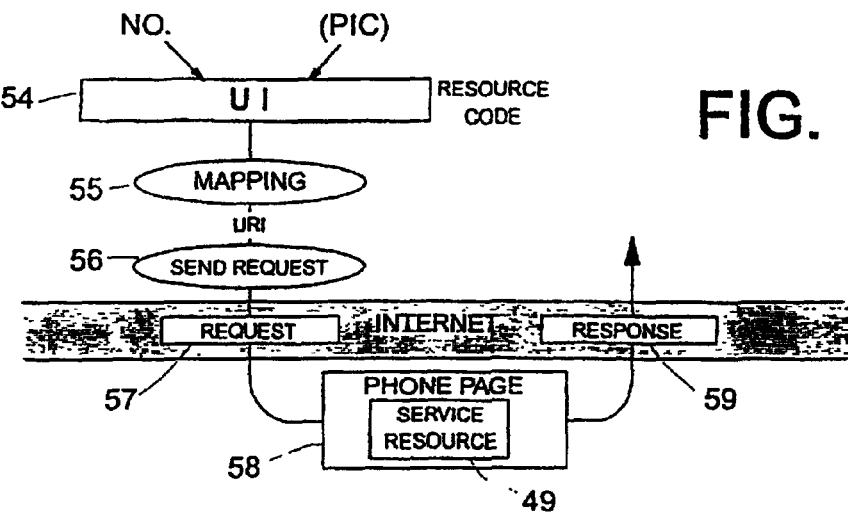


FIG. 9




FIG. 10

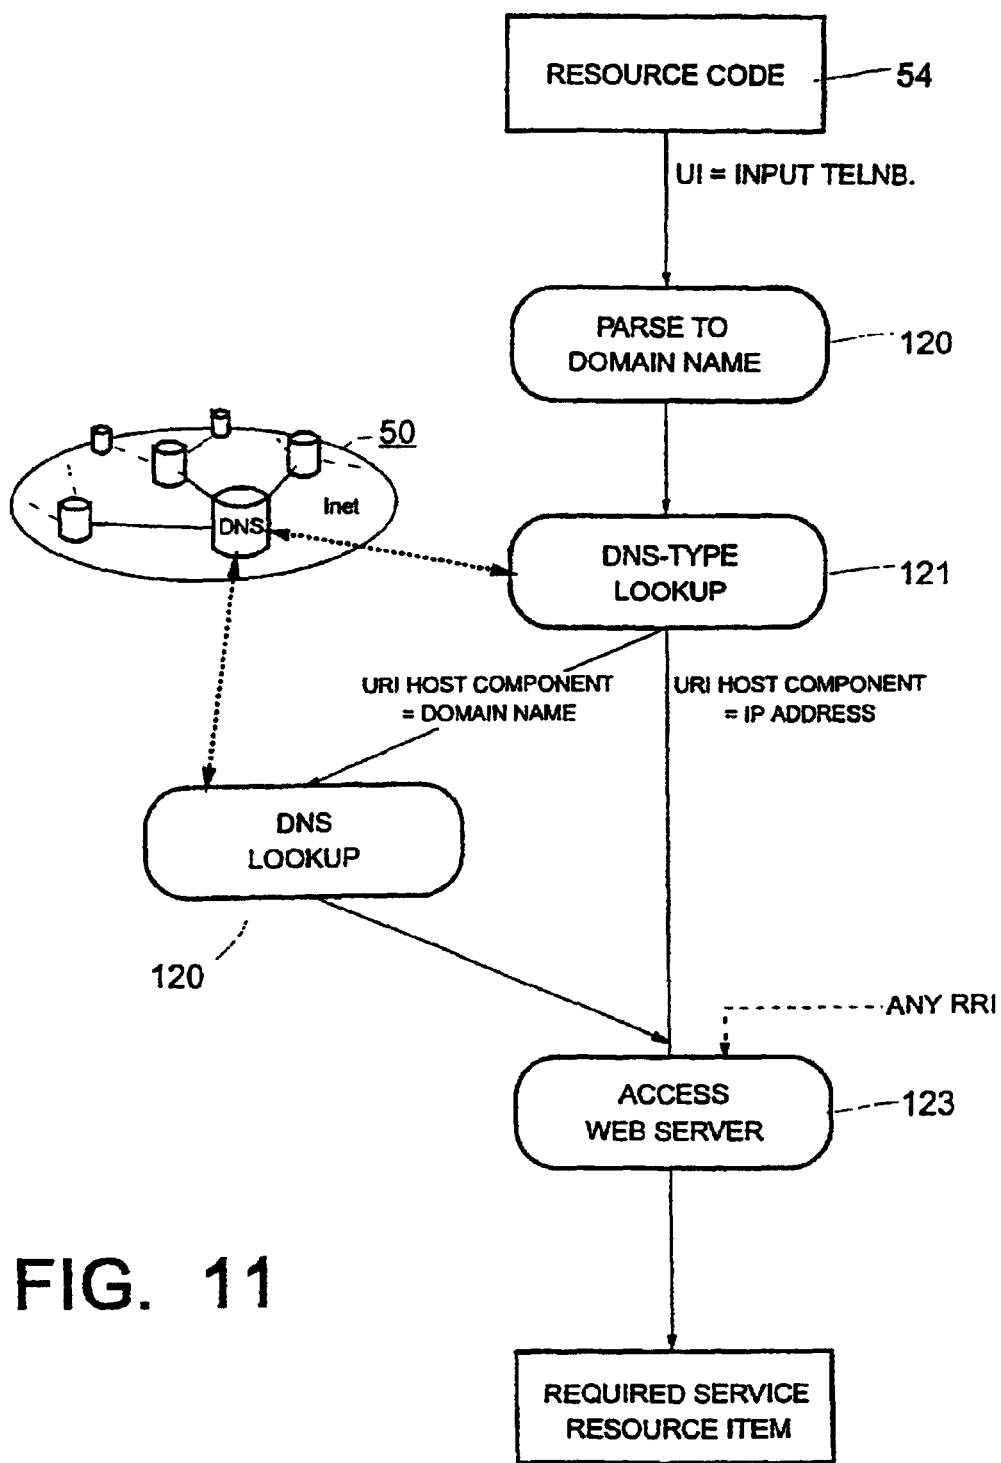


FIG. 11




FIG. 12A




FIG. 12B




FIG. 12C

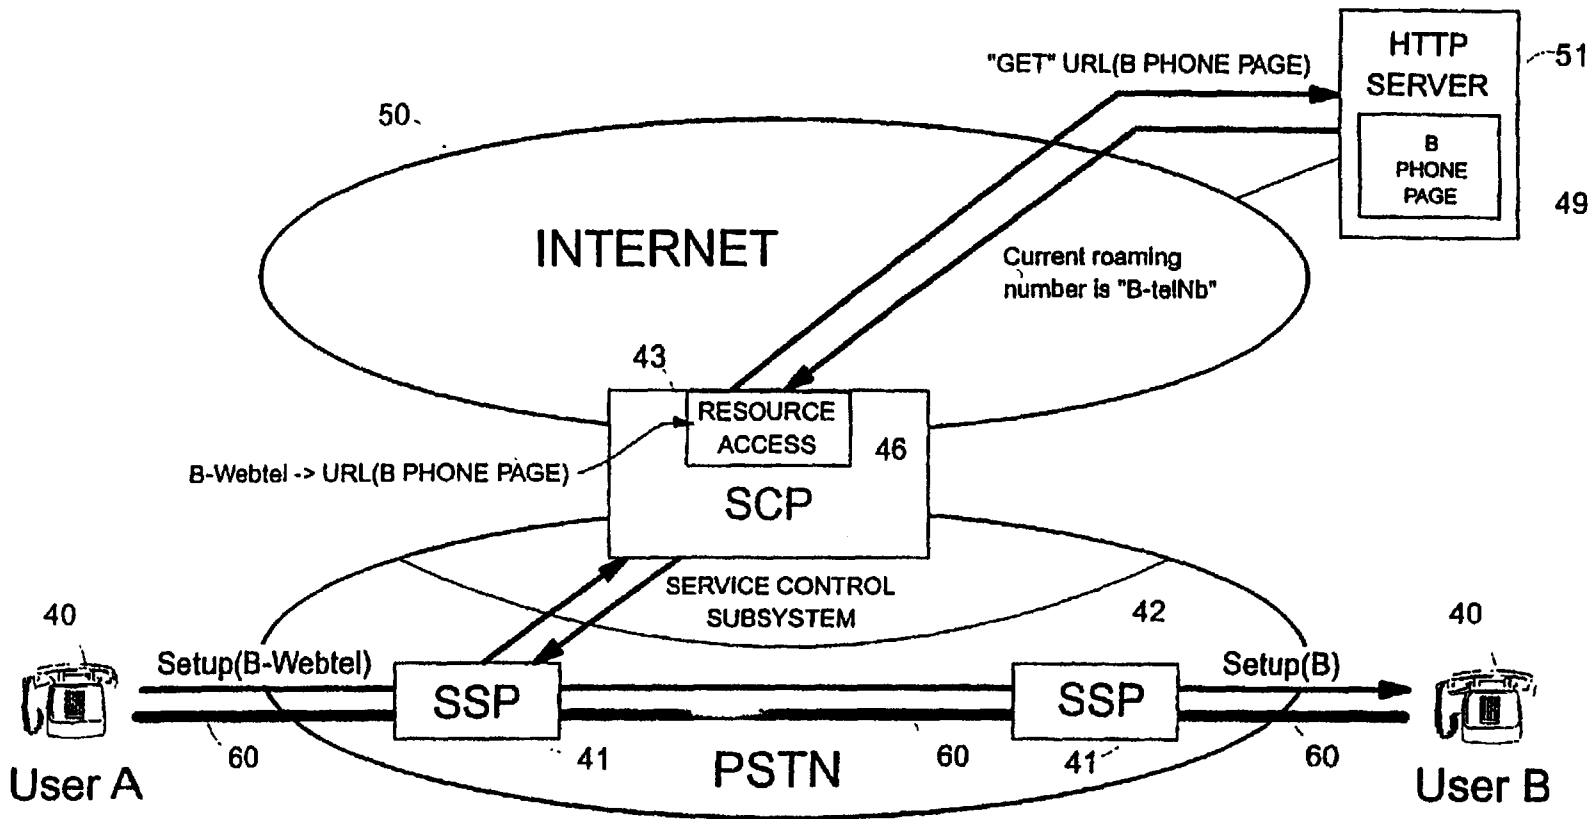


FIG. 13

A82

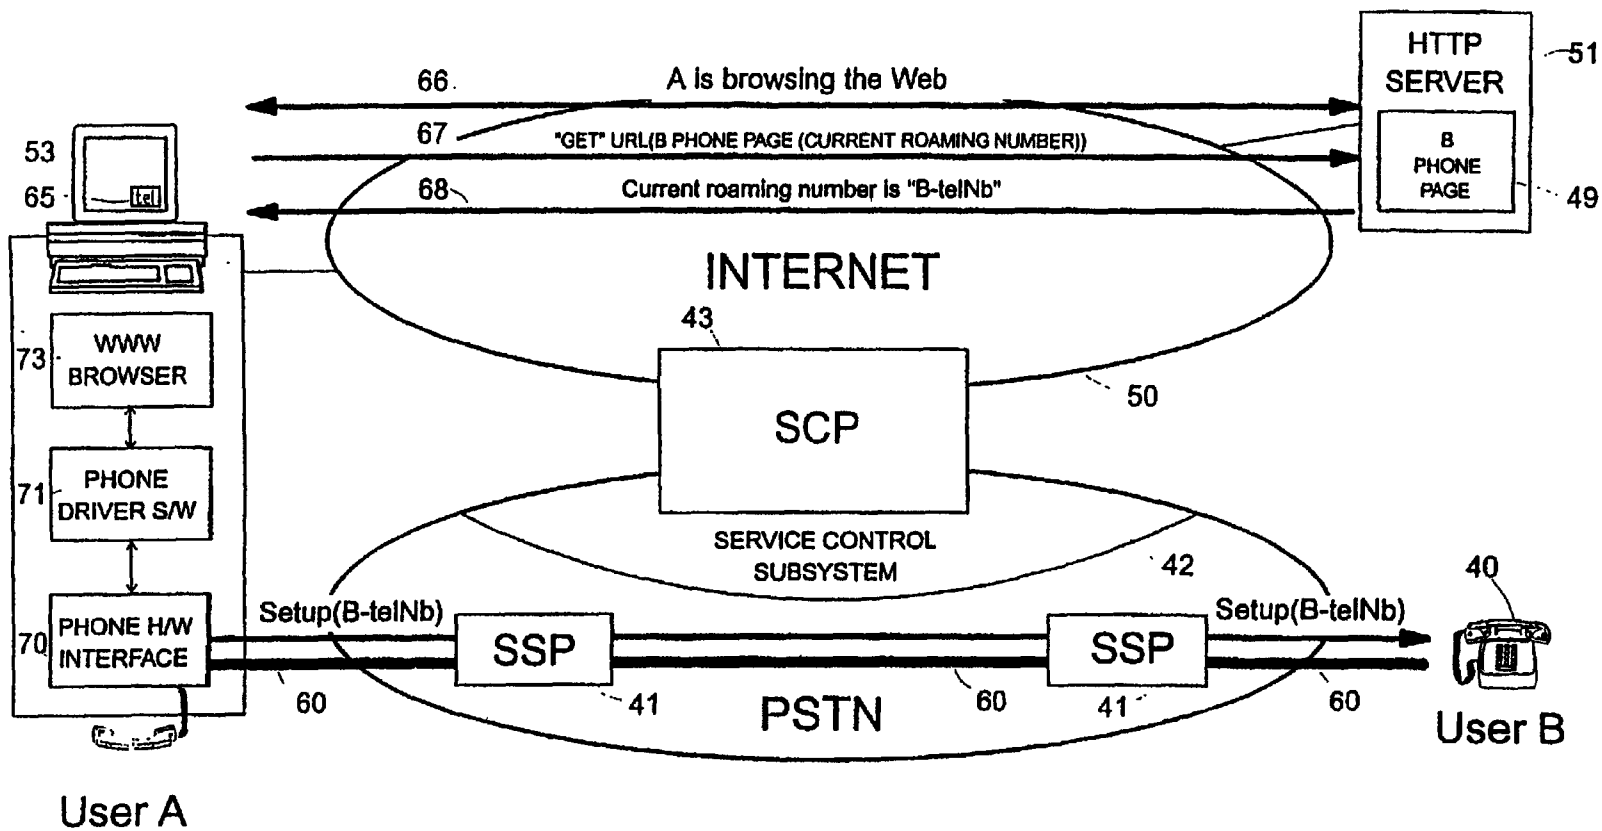


FIG. 14

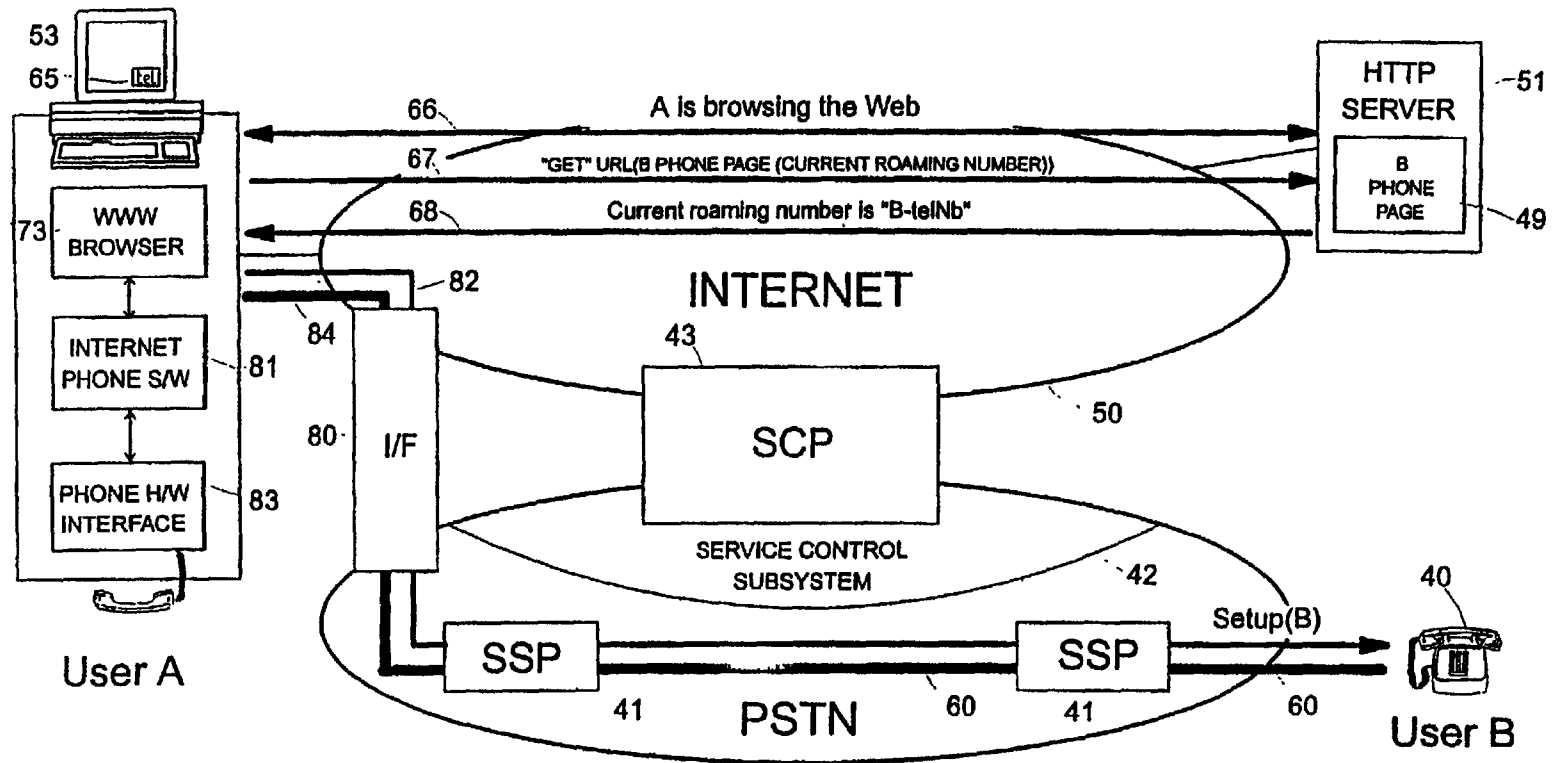


FIG. 15

A84

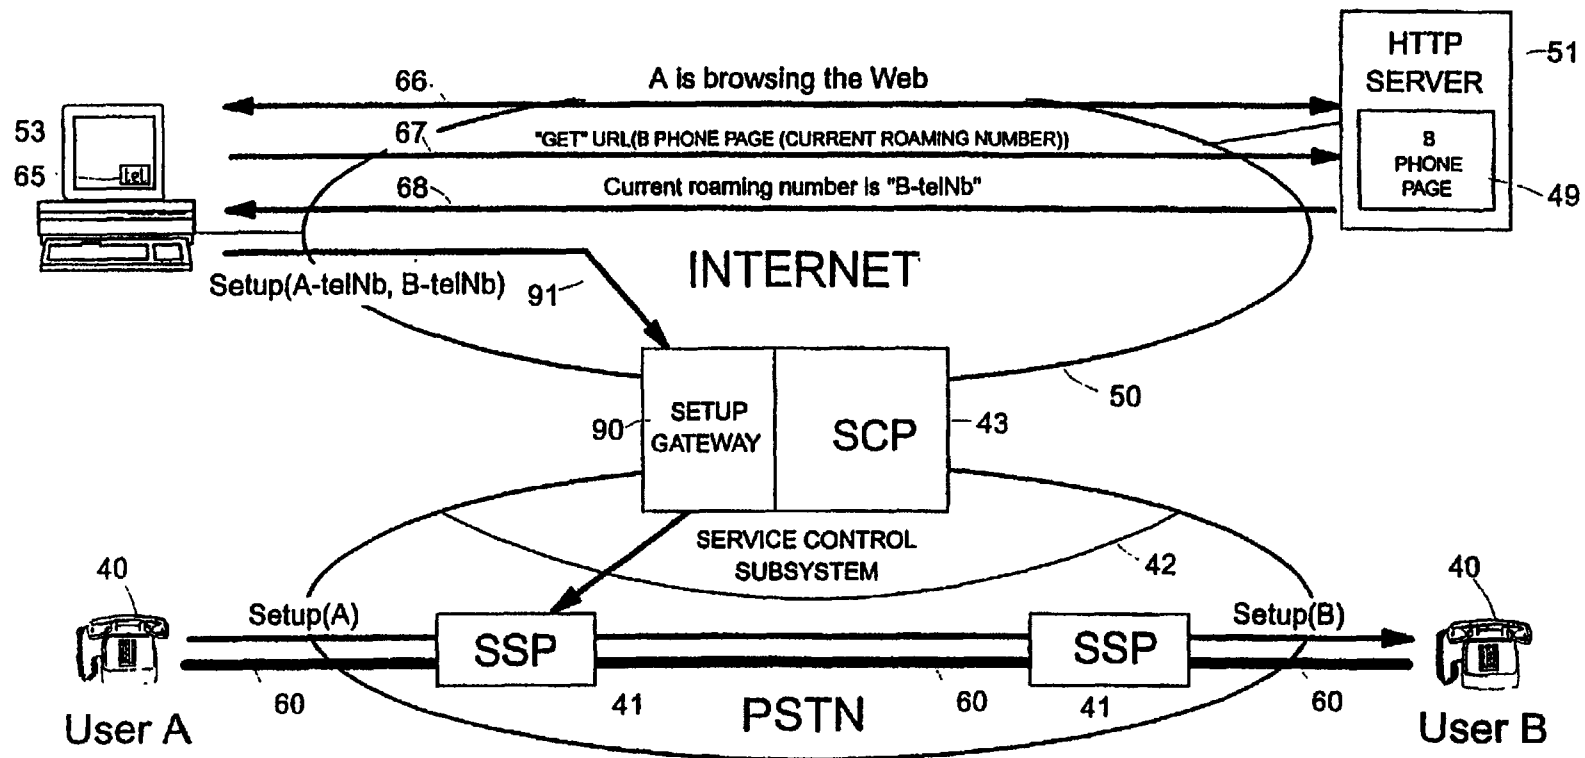


FIG. 16

A85

U.S. Patent

Mar. 14, 2006

Sheet 13 of 14

US 7,012,916 B2

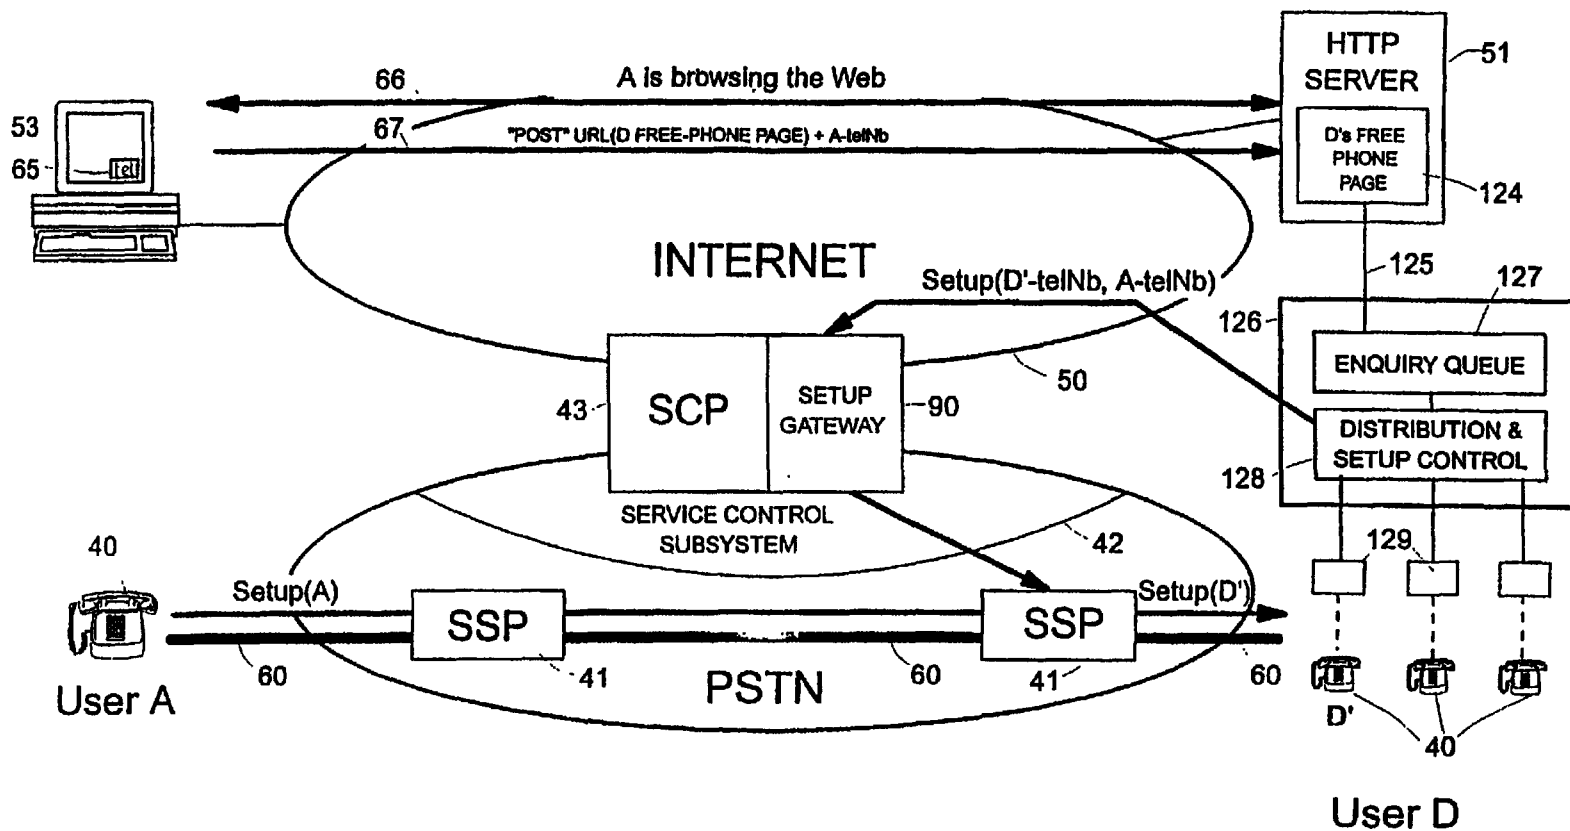


FIG. 17

U.S. Patent

Mar. 14, 2006

Sheet 14 of 14

US 7,012,916 B2




FIG. 18

A87

US 7,012,916 B2

1

METHOD AND APPARATUS FOR ACCESSING COMMUNICATION DATA RELEVANT TO A TARGET ENTITY IDENTIFIED BY A NUMBER STRING

This is a divisional of U.S. Ser. No. 09/077,795, filed on Jun. 5, 1998 now U.S. Pat. No. 6,466,570, which is the U.S. National Stage Application of PCT International Application No. PCT/GB96/03055, filed on Dec. 11, 1996. The present invention relates to methods and apparatus for accessing communication data relevant to a target entity identified by a number string.

FIELD OF THE INVENTION

As used herein, the term "switched telecommunication system" means a system comprising a bearer network with switches for setting up a bearer channel through the network. The term "switched telecommunication system" is to be taken to include not only the existing public and private telephone systems (whether using analogue phones or ISDN-based), but also broadband (ATM) and other switch-based bearer networks that are currently being implemented or may emerge in the future. For convenience, the term "switched telecommunication system" is sometimes shortened herein to telecommunication system.

Reference to a "call" in the context of a switched telecommunication system is to be understood as meaning a communication through a bearer channel set up across the bearer network, whilst references to call setup, maintenance and takedown are to be taken to mean the processes of setting up, maintaining and taking down a bearer channel through the bearer network. Terms such as "call processing" and "call handling" are to be similarly interpreted.

The term "communication system" when used herein should be understood as having a broader meaning than switched telecommunication system, and is intended to include datagram-based communication systems where each data packet is independently routed through a bearer network without following a predetermined bearer channel.

BACKGROUND OF THE INVENTION

Telecommunication companies running PSTNs (Public Switched Telephone Networks) and PLMNs (Public Land Mobile Networks) are in the business of providing communication services and in doing so are providing increasing built-in intelligence in the form of "IN services" such as 800 number services and call forwarding. In contrast, the World Wide Web (WWW), which has seen explosive growth in recent times, is an example of an Internet-based global network providing complex information services. These two worlds, that of the large communications utilities and that of the highly dynamic, pioneer-spirit WWW information culture, are uneasy companions and each plans to encroach on the domain previously occupied by the others; thus telephony services will be offered over the WWW and information services over the public communication infrastructure.

The present invention proposes technologies for a more synergetic relationship between these two worlds than is currently envisaged and in order to place the present invention in context, a review will first be given of each of these two worlds.

Telephone Networks with in Services

The Basic PSTN. The basic service provided by a PSTN (Public Switched Telephone Network) is the interconnection of two telephones (that is, setting up a bearer channel

2

between the telephones) according to a called-party telephone number input at the calling-party telephone. FIG. 1 is a simplified representation of a PSTN providing such a service. In particular, customer premises equipment, CPE, 10 (such as standard analogue telephones, but also more recently ISDN terminals) are connected through an access network 11 to switching points, SPs 12. The SPs 12 form nodes in an inter-exchange network 13 made up of inter-connecting trunks 14 and SPs that are controlled by control entities 15 in the SPs. The control effected by the control entities 15 is determined by signalling inputs received from the CPEs and other SPs, and involves call setup, maintenance and clearance to provide the desired bearer channel between calling CPE and called CPE. Conceptually, the PSTN may be thought of as a bearer network and a control (signalling) network, the function of the latter being to effect call control through the bearer network, namely the control of setup, maintenance and take down of bearer channels through the bearer network; in practice, the bearer and signalling networks may use the same physical circuits and even the same logical channels.

Thus, where the CPE is a traditional dumb telephone, control signalling between the CPE and its local SP is in-band signalling, that is, the signalling is carried on the same channel as used for voice; this signalling is interpreted and converted at the SPs 12 into signalling between SPs that uses a dedicated common-channel signalling network 16 (implemented nowadays using the SS7 protocol suite). Where the CPE is an ISDN terminal, signalling is carried in a separate channel directly from the CPE on an end-to-end. Modern SPs use the ISUP (ISDN User Part) SS7 protocol for inter-exchange call control signalling whether the CPE is a standard telephone or an ISDN terminal.

Telephone Numbering Plans—As certain aspects of the present invention are influenced by the structuring of telephone numbers, a brief description will now be given of the structuring of such numbers. Telephone numbers form an international, hierarchical addressing scheme based on groups of decimal digits. The top level of the hierarchy is administered by the ITU-T, which has allocated single-digit numeric codes to the major geographic zones (for example "1" for North America, "2" for Africa, "3" for Europe, "4" for Europe, "5" for South America and Cuba, etc.). Within each zone countries are assigned 2 or 3 digit codes, so that within zone 3 France is "33", and within zone 4 the UK is "44". Administration of the numbering plan within a country is delegated to a national body, such as the Office of Telecommunications ("OfTel") in the UK. The following further description is based on the UK numbering plan, but the scheme described will be recognised as having widespread applicability.

In the UK all national numbers are prefixed by a code from 01 to 09 (the '0' prefix is dropped in international dialling). The currently assigned codes are "01" for Geographic Area Codes, "02" for Additional Geographic Area Codes, "04" for Mobile Services, "07" for Personal Numbers, and "08" for Special Service (freephone, information). Normal wireline PSTN subscriber telephone numbers are allocated from the Geographic Area Code codes, and currently only codes prefixed by 01 are allocated. Geographic area codes are presently 3 or 4 digits (excluding the leading '0') and there are currently 638 geographic areas each with its own code. A full national UK dialled number takes two forms:

0	171	634 8700
	area code	local number (7 digit)
0	1447	456 987
	area code	local number (6 digit)

The first case has the '0' prefix, a 3 digit area code and a 7 digit local number, and the second case has the '0' prefix, a 4 digit area code, and a 6 digit local number. Further interpretation of the local number will take place within the area exchange, as even a 6 digit address space is too large for a single switch, and for a typical local area several switches may be needed to host the required number of subscriber lines. This interpretation is opaque and is a matter for the area service provider.

In the current PSTN the inherently hierarchical and geographic interpretation of telephone numbers is mirrored by the physical architecture of the network. A telephone number is structured in a way that makes it easy to route a call through the network. At each step, the prefix of the number provides information about the current routing step, and the suffix (perhaps opaquely) provides information about subsequent routing steps; as long as a switch knows how to parse a prefix and carry out a routing step, it does not need to understand the content of the suffix, which is left for subsequent routing steps. For this reason the international and national switching fabric is also organised hierarchically.

Intelligent Networks. Returning now to a consideration of the current telephone network infrastructure, in addition to basic call handling, an SP may also serve to provide what are called IN (Intelligent Network) services; in this case the SP is termed a service switching point, SSP. An SSP 25 is arranged to suspend call processing at defined points-in-call upon particular criteria being met, and to delegate the continuation of call processing to a service control subsystem providing a service control function (SCF) either in the form of a service control point, SCP 17 (see FIG. 2) or an Adjunct 18. The Adjunct 18 is directly associated with an SSP 25 whilst the SCP 17 and SSP 25 communicate with each other via an extended common channel signalling (CCS) network 16 that may include signal transfer points (STP) 19. The SCP 17 may be associated with more than one SSP 25. Both the SCP 17 and Adjunct 18 provide a service logic execution environment (SLEE) 20 in which instances of one or more service logic programs (SLP) 21 can execute. The SLEE 20 and SLP 21 together provide service control functionality for providing services to the SSP 25.

Service logic running in an SCP or Adjunct will generally make use of subscriber information stored in a service data function (SDF) 22 that may be integral with the SCP/Adjunct or partially or wholly separate therefrom. The service data function (SDF), like the service control function (SCF) forms part of the service control subsystem of the PSTN. It may be noted that some or all of the service control function may be built into the PSTN switches themselves.

In addition to the SCP 17 and Adjunct 18, the FIG. 2 network includes an intelligent peripheral (IP) 23. The IP 23 provides resources to the SSP 25 such as voice announcements and DTMF digit collection capabilities. The network will also include an operation system (not shown) that has a general view of the network and its services and performs functions such as network monitoring and control.

In operation, when the SSP 25 receives a call, it examines internal trigger conditions and, possibly, user information (eg dialled digits) to ascertain if the call requires a service to

be provided by the service control subsystem 17, 18; the checking of trigger conditions may be carried out at several different points in call processing. Where the SSP 25 determines that a service is required it messages the service control subsystem (either SCP 17 or Adjunct 18) requesting the desired service and sending it a logic representation of the call in terms of its connectivity and call processing status. The service control subsystem then provides the requested service and this may involve either a single interaction between the SSP and service control subsystem or a session of interactions. A typical service is call forwarding which is a called-party service giving expression to an end-user requirement as simple as "if you call me on number X and it rings ten times, try calling number Y". In this case, it is the SSP local to the called end-user that triggers its associated SCP (or Adjunct) to provide this service; it will, of course, be appreciated that the SSP must be primed to know that the service is to be provided for a called number X.

The above-described model for the provision of IN services in a PSTN can also be mapped onto PLMNs (Public Land Mobile Networks) such as GSM and other mobile networks. Control signalling in the case of a mobile subscriber is more complex because in addition to all the usual signalling requirements, there is also a need to establish where a call to a mobile subscriber should be routed; however, this is not a very different problem from a number of called-party IN services in the PSTN. Thus in GSM, the service-data function (SDF) is largely located in a system named a Home Location Register (HIR) and the service control function in a system named a Visitor Location Register (VLR) that is generally associated on a one-to-one basis with each SSP (which in GSM terminology is called a Mobile Switching Centre, MSC).

Because subscribers are mobile, the subscriber profile is transported from the HLR to whichever VLR happens to be functionally closest to be mobile subscriber, and from there the VLR operates the (fixed) service using the subscriber profile and interacts with the SSP. The HLR and VIR thus constitute a service control subsystem similar to an SCP or Adjunct with their associated databases.

It is, of course, also possible to provide IN services in private telephone systems and, in this case, the service control function and service data function are generally either integrated into a PABX (Private Automatic Branch Exchange) or provided by a local computer. The service control subsystem, whilst present, may thus not be a physically distinct from the PABX.

The above-described general architectural framework for providing IN services has both strengths and flaws. Its main strength is that it works and many services have been successfully deployed, such as 800 number services, credit card calling, voicemail, and various call waiting and redirection services. However, despite years of standardisation, services are still implemented one-at-a-time on proprietary platforms and do not scale well. The approach has been based on large, fault-tolerant systems which provide services for hundreds of thousands or even millions of subscribers and take years to deploy. Furthermore, since the networks used to support these services also constitute the basic telephone infrastructure, anything attached to these networks must be rigorously vetted. Additionally, each country and operator tends to have local variations of the so-called standards making it difficult to supply standard products and thereby braking the dynamics of competition. The World Wide Web

In contrast to the slow deliberate progress of the telephone infrastructure, the WWW has grown explosively from its

US 7,012,916 B2

5

inception in 1989 to become the primary electronic information distribution service in terms of spread, availability and richness of information content. Anyone can, for a modest outlay, become an information provider with a world-wide audience in a highly interconnected information architecture.

The WWW is a client-server application running over the Internet and using a client-server protocol which mandates only the simplest of exchanges between client and server. This protocol is HTTP (Hyper Text Transfer Protocol) which is optimised for use over TCP/IP networks such as the Internet; the HTTP protocol can, however, also be used over networks using different communication protocol stacks.

Since the availability of literature concerning the WWW has seen the same sort of growth as the WWW itself, a detailed description of the WWW, HTTP and the Internet will not be given herein. An outline description will, however, be given with attention being paid to certain features of relevance to the present invention.

The WWW uses the Internet for interconnectivity. Internet is a system that connects together networks on a world-wide basis. Internet is based on the TCP/IP protocol suite and provides connectivity to networks that also use TCP/IP. For an entity to have a presence on the Internet, it needs both access to a network connected to the Internet and an IP address. IP addresses are hierarchically structured. Generally an entity will be identified at the user level by a name that can be resolved into the corresponding IP address by the Domain Name System (DNS) of the Internet. Because the DNS or adaptations of it are fundamental to at least certain embodiments of the invention described hereinafter, a description will next be given of the general form and operation of the DNS.

The Domain Name System—The DNS is a global, distributed, database, and without its performance, resilience and scalability much of the Internet would not exist in its current form. The DNS, in response to a client request, serves to associate an Internet host domain name with one or more Registration Records (RR) of differing types, the most common being an address (A) record (such as 15.144.8.69) and mail exchanger (MX) records (used to identify a domain host configured to accept electronic mail for a domain). The RRs are distributed across DNS name servers world-wide, these servers cooperating to provide the domain name translation service; no single DNS server contains more than a small part of the global database, but each server knows how to locate DNS servers which are “closer” to the data than it is. For present purposes, the main characteristics of the DNS of interest are:

The host name space is organised as a tree-structured hierarchy of nodes with each host having a corresponding leaf node; each node has a label (except the root node) and each label begins with an alphabetic character and is followed by a sequence of alphabetic characters or digits. The full, or “fully qualified” name of a host is the string of node labels, each separated by a “.”, from the corresponding leaf node to the root node of the hierarchy, this latter being represented by a terminating “.” in the name. Thus a host machine “fred” of Hewlett-Packard Laboratories in Bristol, England will have a fully qualified domain name of “fred.hpl.hp.com.” (note that if a host name does not have a terminal “.” it is interpreted relative to the current node of the naming hierarchy).

Each host has one or more associated Registration Records (RRs).

There are a plurality of DNS servers each with responsibility for a subtree of the name space. A DNS server

6

will hold RRs for all or part of its subtree—in the latter case it delegates responsibility for the remainder of the subtree to one or more further DNS servers. A DNS server knows the address of any server to which it has delegated responsibility and also the address of the server which has given it the responsibility for the subtree it manages. The DNS servers thus point to each other in a structuring reflecting that of the naming hierarchy.

An application wishing to make use of the DNS does so through an associated “resolver” that knows the address of at least one DNS server. When a DNS server is asked by this resolver for an RR of a specified host, it will return either the requested RR or the address of a DNS server closer to the server holding the RR in terms of traversal of the naming hierarchy. In effect, the hierarchy of the servers is ascended until a server is reached that also has responsibility for the domain name to be resolved; thereafter, the DNS server hierarchy is descended down to the server holding the RR for the domain name to be resolved.

The DNS uses a predetermined message format (in fact, it is the same for query and response) and uses the IP protocols.

These characteristics of the DNS may be considered as defining a “DNS-type” system always allowing for minor variations such as in label syntax, how the labels are combined (ordering, separators), the message format details, evolutions of the IP protocols etc.

Due to the hierarchical naming structure, it is possible to delegate responsibility for administering domains (subtrees) of the name space recursively. Thus, the top-level domains are administered by InterNic (these top-level domains include the familiar ‘com’, ‘edu’, ‘org’, ‘int’, ‘net’, ‘mil’ domains as well as top-level country domains specified by standard two-letter codes such as ‘us’, ‘uk’, ‘fr’ etc.). At the next level, by way of example Hewlett-Packard Company is responsible for all names ending in ‘hp.com’ and British Universities are collectively responsible for all names ending in ‘ac.uk’. Descending further, and again by way of example, administration of the domain ‘hpl.hp.com’ is the responsibility of Hewlett-Packard Laboratories and administration of the subtree (domain) ‘newcastle.ac.uk’ is the responsibility of the University of Newcastle-upon-Tyne.

FIG. 3 illustrates the progress of an example query made from within Hewlett-Packard Laboratories. The host domain name to be resolved is ‘xy.newcastle.ac.uk’, a hypothetical machine at the University of Newcastle, United Kingdom. The query is presented to the DNS server responsible for the “hpl.hp.com” subtree. This server does not hold the requested RR and so responds with the address of the “hp.com” DNS server; this server is then queried and responds with the address of the ‘com’ DNS server which in turn responds with the address of the ‘.’ (root) DNS server. The query then proceeds iteratively down the ‘uk’ branch until the ‘newcastle.ac.uk’ server responds with the RR record for the name ‘xy’ in its subtree.

This looks extremely inefficient, but DNS servers are designed to build a dynamic cache, and are initialised with the addresses of several root servers, so in practice most of the iterative queries never take place. In this case the ‘hpl.hp.com’ DNS server will know the addresses of several root servers, and will likely have the addresses of ‘uk’ and ‘ac.uk’ servers in its cache. The first query to the ‘hpl.hp.com’ server will return the address of the ‘ac.uk’ server. The second query to the ‘ac.uk’ server will return the address of the ‘newcastle.ac.uk’ server, and the third query will return

the RR in question. Any future queries with a ‘newcastle.ac.uk’ prefix will go direct to the newcastle DNS server as that address will be retained in the “hpl.hp.com” DNS server cache. In practice names within a local subtree are resolved in a single query, and names outside the local subtree are resolved in two or three queries.

Rather than a resolver being responsible for carrying out the series of query iterations required to resolve a domain name, the resolver may specify its first query to be recursive in which case the receiving DNS server is responsible for resolving the query (if it cannot directly return the requested RR, it will itself issue a recursive query to a ‘closer’ DNS server, and so on).

It should also be noted that in practice each DNS server will be replicated, that is, organised as a primary and one or more secondaries. A primary DNS nameserver initialises itself from a database maintained on a local file system, while a secondary initialises itself by transferring information from a primary. A subtree will normally have one primary nameserver and anything up to ten secondaries—the limitation tends to be the time required by the secondaries to update their databases from the primary. The primary database is the master source of subtree information and is maintained by the domain DNS administrator. The secondaries are not simply standby secondaries but each actively participates in the DNS with dependent servers that point to it rather than to the corresponding primary.

DNS implementations, such as BIND, are widely available as a standard part of most UNIX systems, and can claim to be among the most robust and widely used distributed applications in existence.

Operation of the WWW Referring now to FIG. 4 of the accompanying drawings, access to the Internet 30 may be by direct connection to a network that is itself directly or indirectly connected to the Internet; such an arrangement is represented by terminal 31 in FIG. 4 (this terminal may, for example, be a Unix workstation or a PC). Having a connection to the Internet of this form is known as having ‘network access’. Any entity that has network access to the Internet may act as a server on the Internet provided it has sufficient associated functionality; in FIG. 4, entity 32 with file store 37 acts as a server.

Many users of the WWW do not have network access to the Internet but instead access the Internet via an Internet service provider, ISP, 33 that does have network access. In this case, the user terminal 34 will generally communicate with the ISP 33 over the public telephone system using a modem and employing either SLIP (Serial Line Interface Protocol) or PPP (Point-to-Point Protocol). These protocols allow Internet packets to traverse ordinary telephone lines. Access to the Internet of this form is known as “dialup IP” access. With this access method, the user terminal 34 is temporarily allocated an IP address during each user session; however, since this IP address may differ between sessions, it is not practical for the entity 34 to act as a server.

A cornerstone of the WWW is its ability to address particular information resources by means of an Uniform Resource Identifier (URI) that will generally be either a Uniform Resource Locator (URL) that identifies a resource by location, or a Uniform Resource Name (URN) that can be resolved into an URL. By way of example, a full or “absolute” URL will comprise the following elements:

scheme	this is the access scheme to be used to access the resource of interest;
host	the Internet host domain name or IP address;
port	the host port for the (TCP) connection;
abs-path	the absolute path of the resource on the host.

In fact, the ‘port’ may be omitted in which case port 80 is assumed.

FIG. 5 of the accompanying drawings shows an example URL for the Hewlett-Packard products welcome page. In this case, the elements are:

scheme	http
host	www.hp.com
port	omitted (port 80 assumed)
abs-path	Products.html

The HTTP protocol is based on a request/response paradigm. Referring again to FIG. 4 of the drawings, given a particular URI identifying a resource 30 to be accessed, a client establishes a connection with the server 31 corresponding to the “host” element of the URI and sends a request to the server. This request includes a request method, and the “Request-URI” (which is generally just the absolute path of the resource on the server as identified by the “abs-path” element of the URI); the request may include additional data elements. The server 31 then accesses the resource 36 (here held on storage 37) and responds and this response may include an entity of a type identified by a MIME (Multipurpose Internet Mail Extensions) type also included in the response.

The two main request methods are:

GET—This method results in the retrieval of whatever information (in the form of an entity) is identified by the Request-URI. It is important to note that if the Request-URI refers to a data-producing process, it is the produced data which is returned as the entity in the response and not the source text of the process.

POST—This method is used to request that the destination server accept the entity enclosed in the request as a new subordinate of the resource identified by the Request-URI. The POST method can be used for annotation of existing resources, providing a message to a bulletin board, providing data to a data-handling process (for example, data produced as the result of submitting a form), and extending a database through an append operation.

In summary, the GET method can be used to directly retrieve data, or to trigger any process that will return an entity (which may either be data or a simply an indication of the result of running the process). The POST method is used for registering data and specifying this method is also effective to trigger a process in the server to handle the posted data appropriately.

The passing of information to a process triggered to run on a server using either the GET or POST method is currently done according to an interface called the Common Gateway Interface (CGI). The receiving process is often written in a scripting language though this is not essential. Typically, the triggered server script is used for interfacing to a database to service a query included in a GET request. Another use, already referred to, is to append data associated with a POST request to a database.

US 7,012,916 B2

9

Other important factors in the success of the WWW is the use of the HyperText Markup Language (HTML) for representing the makeup of documents transferred over the WWW, and the availability of powerful graphical Web browsers, such as Netscape and Mosaic, for interpreting such documents in a client terminal to present them to a user. Basically, HTML is used to identify each part of a document, such as a title, or a graphic, and it is then up to the browser running in the client terminal to decide how to display each document part. However, HTML is more than this—it also enables a URI and a request method to be associated with any element of a document (such as a particular word or an image) so that when a user points to and clicks on that element, the resource identified by the URI is accessed according to the scheme (protocol) and request method specified. This arrangement provides a hyperlink from one document to another. Using such hyperlinks, a user at a client terminal can skip effortlessly from one document downloaded from a server on one side of the world, to another document located on a server on the other side of the world. Since a document created by one author may include a hyperlink to a document created by another, an extremely powerful document cross-referring system results with no central bureaucratic control.

Hyperlinks are not the only intelligence that can be built into an HTML document. Another powerful feature is the ability to fill in a downloaded “Form” document on screen and then activate a ‘commit’ graphical button in order to have the entered information passed to a resource (such as a database) designed to collect such information. This is achieved by associating the POST request method with the ‘commit’ button together with the URI of the database resource; activating the ‘commit’ button results in the entered information being posted to the identified resource where it is appropriately handled.

Another powerful possibility is the association of program code (generally scripts to be interpreted) with particular documents elements, such as graphical buttons, this code being executed upon the button being activated. This opens up the possibility of users downloading program code from a resource and then running the code.

It will be appreciated by persons skilled in the art that HTML is only one of several currently available scripting languages delivering the functionality outlined above and it may be expected that any serious Web browser will have built-in support for multiple scripting languages. For example, Netscape 2.0 supports HTML 3.0, Java and LiveScript (the latter being Netscape proprietary scripting Language).

The importance of the role of the graphical Web browser itself should not be overlooked. As well as the ability to support multiple scripting languages, a Web browser should provide built-in support for standard media types, and the ability to load and execute programs in the client, amongst other features. These browsers may be viewed as operating systems for WWW interaction.

WWW and the Telephone Network

It is possible to provide a telephony service over the Internet between connected terminals by digitising voice input and sending it over the Internet in discrete packets for reassembly at the receiving terminal. This is an example of a communication service on the Internet. Conversely, it is possible to point to a variety of information services provided over the telephone system, such as the Minitel system widely available in France. However, these encroachments into each others traditional territories pose no real threat to either the Internet or the public telephone system.

10

Of more interest are areas of cooperative use of the Internet and the telephone system. In fact, one such area has existed for some considerable time and has been outlined above with reference to FIG. 4, namely the use of a modem link over the PSTN from a user computer 34 to an Internet service provider 33 in order to obtain dialup IP access to the Internet. This cooperative use is of a very simple nature, namely the setting up of a bearer channel over the PSTN for subsequently generated Internet traffic; there is no true interaction between the Internet and the PSTN.

Another known example of the cooperative use of the Internet and PSTN is a recently launched service by which an Internet user with a sound card in his/her terminal computer can make a voice call to a standard telephone anywhere in the world. This is achieved by transferring digitised voice over the Internet to a service provider near the destination telephone; this service provider then connects into the local PSTN to access the desired phone and transfers across into the local PSTN the voice traffic received over the Internet. Voice input from the called telephone is handled in the reverse manner. Key to this service is the ability to identify the service provider local (in telephony charging terms) to the destination phone. This arrangement, whilst offering the prospect of competition for the telecom operators for long distance calls, is again a simple chaining together of the Internet and PSTN. It may, however, be noted that in this case it is necessary to provide at least a minimum of feedback to the Internet calling party on the progress of call set to the destination telephone over the PSTN local to that telephone; this feedback need only be in terms of whether or not the call has succeeded.

According to the present invention, there is provided a method of accessing communication data relevant to a target entity identified by a number string, said method comprising the steps of:

- (a)—storing in the domain name system (DNS) of the Internet records each associated with a corresponding domain name and holding an URI for locating communications data associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a substantial portion of the number string into at least a part of said domain name;
- (b)—applying said process to the said number string identifying the target entity whereby to form the related domain name;
- (c)—supplying the domain name formed in step (b) to the DNS to retrieve the URI held in the corresponding said record; and
- (d)—using the URI retrieved in step (c) to access said communication data.

According to another aspect of the present invention, there is provided a method of accessing communication data relevant to a target entity identified by a number string, said method comprising the steps of:

- (a)—storing in the domain name system (DNS) of the Internet records each associated with a corresponding domain name and holding an at least part-formed URL, including access scheme and host name, of an item of communications data, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a substantial portion of the number string into at least a part of said domain name;
- (b)—applying said process to the said number string identifying the target entity whereby to form the related domain name; and

US 7,012,916 B2

11

- (c)—supplying the domain name formed in step (b) to the DNS to retrieve the at least part-formed URL held in the corresponding said record;
- (d)—using the at least part-formed URL retrieved in step (c) to access said communication data.

According to a further aspect of the present invention, there is provided a method of discovering communications endpoint address data for contacting a target entity identified by a number string, said method comprising the steps of:

- (a)—storing in the domain name system (DNS) of the Internet records each associated with a corresponding domain name and holding an URL of a resource that has access to multiple items of communications endpoint address data, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a substantial portion of the number string into at least a part of said domain name, the number strings being in telephone-number form;
- (b)—applying said process to the said number string identifying the target entity whereby to form the related domain name; and
- (c)—supplying the domain name formed in step (b) to the DNS to retrieve the URL held in the corresponding said record;
- (d)—using the URL to access the corresponding said resource and supply it with an indicator of the desired item of communications endpoint address data, this data then being returned by the resource.

According to a still further aspect of the present invention, there is provided a method of accessing communication data relevant to a target entity identified by a number string, said method comprising the steps of:

- (a)—storing in a DNS-type database system, records each associated with a corresponding domain name and holding an URI for locating communications data associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a substantial portion of the number string into at least a part of said domain name;
- (b)—applying said process to the said number string identifying the target entity whereby to form the related domain name; and
- (c)—supplying the domain name formed in step (b) to the DNS-type database system to retrieve the URI held in the corresponding said record;
- (d)—using the URI retrieved in step (c) to access said communication data.

The present invention also encompasses clients-focused methods corresponding to the above overall methods of the invention, and apparatus for implementing aspects of the methods of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will now be described, by way of non-limiting example, with reference to the accompanying diagrammatic drawings, in which:

FIG. 1 is a simplified diagram of a standard PSTN;

FIG. 2 is a simplified diagram of a known PSTN with IN service capability;

FIG. 3 is a diagram illustrating host domain name resolution by the DNS of the Internet;

FIG. 4 is a diagram illustrating the functioning of the World Wide Web;

12

FIG. 5 is a diagram illustrating the format of a standard URL;

FIG. 6 is a diagram of a first arrangement in which service resource items are held on HTTP servers accessible both by the service control subsystem of a PSTN and by Web users;

FIG. 7 is a diagram illustrating the processing of a service request by the SCP of FIG. 6;

FIG. 8 is a diagram illustrating the format of a resource code used by the FIG. 6 SCP when accessing a service resource item;

FIG. 9 is a diagram illustrating the process of accessing a service resource in the case where the service code does not include an RRI part;

FIG. 10 is a diagram illustrating the process of accessing a service resource in the case where the service code includes an RRI part;

FIG. 11 is a diagram illustrating the derivation of the URI of a service resource by parsing an input telephone number in accordance with the present invention;

FIG. 12A is a diagram depicting a name space (the "telname space") constituted by the domain names derived by a parsing of a predetermined set of telephone numbers;

FIG. 12B is a diagram depicting the incorporation of the telname space without fragmentation into the DNS;

FIG. 12C is a diagram depicting the incorporation of the telname space in fragmented form into the DNS;

FIG. 13 is a diagram illustrating the overall operation of the FIG. 6 arrangement in providing a roaming number service in response to a telephone number being dialled at a standard phone;

FIG. 14 is a diagram illustrating the overall operation of the FIG. 6 arrangement when utilised by a Web user in setting up a call through a telephone interface integrated into the user's Web terminal;

FIG. 15 is a diagram illustrating the overall operation of an arrangement in which an interface is provided between the PSTN and the Internet for telephone traffic;

FIG. 16 is a diagram illustrating the overall operation of an arrangement in which a call setup gateway is provided between the Internet and the PSTN;

FIG. 17 is a diagram illustrating the overall operation of an arrangement in which a freephone service is implemented for Web users; and

FIG. 18 is a diagram similar to FIG. 6 illustrating the provision of a distributed processing environment for inter-connecting elements of the service control subsystem of the PSTN.

BEST MODE OF CARRYING OUT THE INVENTION

FIG. 6 illustrates an arrangement for the provision of services in a PSTN conventionally comprising an inter-exchange network 13 (including trunks and switches at least some of which are SSPs 41 with associated IPs), an access network 11 connecting customer premise equipment (here shown as telephones 40) to the network 13, and a service control subsystem 42 including at least one SCP for providing services to the SSPs 41 upon request. It will be appreciated that the FIG. 6 representation of a PSTN is highly diagrammatic.

The SCP 43 may operate in a conventional manner responding to service requests from SSPs 41 to run specific service logic on particular data according to information contained in the service request, and to send back to the

US 7,012,916 B2

13

requesting SSP appropriate instructions for effecting call set up. A service request is generated by the SSP in response to predetermined trigger conditions being met at a trigger check point, there being one or more such check points in the course of handling a call (it may be noted that where the trigger conditions have been downloaded to the SSP from the SCP then it could be said that the SSP is responding to an information request by the SCP when contacting the SCP upon the trigger conditions being met—however, in the present specification, this initial communication from the SSP to the SCP will be referred to as a “service request”).

The SCP 43 is also provided with a network access interface 44 to the Internet 50 in order to make use of certain service resource items 49 (also referred to below simply as “service resources”) during the course of processing at least certain service requests from the SSPs 41. These service resources 49 are held as WWW pages on HTTP servers 51 (more particularly, on service resource databases 52 of these servers 51). The WWW pages containing these service resources are referred to below as “phone” pages. The servers 51 are connected to the Internet and the phone pages are read accessible using respective URLs or URNs (for convenience, the more general term URI will be used hereinafter to mean the Internet-resolvable indicator of the location of a phone page).

The service resources may be service logic or service data and may be used by an otherwise standard service logic program running on the SCP, by accessing the phone page of the required resource using the appropriate URI. In certain cases, the service resources 49 may provide substantially all of the service control and data associated with a particular service. In this case, the service logic program running in the SCP 43 is of skeleton form, being instantiated on receipt of a service request and then serving to initiate service resources access and to return the results of this access to the entity that made the service request. In fact, according to this approach, the SCP could be implemented simply as a platform for fetching and executing phone-page service logic and would not need to have the complex provisioning and management systems for such logic as is required by standard SCP platforms; SCPs could then become more ubiquitous, possibly being associated with every SSP.

FIG. 7 is a flow chart illustrating the progress of events in the case where the SCP 43 handles a service request by accessing a phone-page service resource. Upon receipt of a service request in an INAP message (step 100), SCP 43 decodes the TCAP/INAP message structure in standard manner (steps 101 and 102) well understood by persons skilled in the art. Next, SCP 43 instantiates a service logic program, SLP, to handle the request (step 103). This SLP is then responsible for looking up the URL of the required service resource as determined from information contained in the service request (steps 104, 105). For example, if the service request relates to a called-party service, then the required resource will be indicated by the dialled number and the latter will be used to derive the URL of the resource. Once the URL of the desired service resource has been ascertained, a resource request (for example, in the form of an HTTP request message) is sent over the Internet to the corresponding server holding the desired service resource (step 106); a correlation ID is also passed with the resource request to enable a response from the latter to be linked with the appropriate SLP instance. A timer is also started (step 107).

If a response is received from the accessed resource before the expiration of a time-out period (tested in step

14

108), then the response, which is usually in the form of a destination number, is supplied to the appropriate SLP as identified using the correlation ID passed with the response (step 109). An INAP/TCAP response message is then prepared and sent to the entity that made the original service request (steps 110 and 111) after which the SLP instance is terminated (113).

If in step 108, a time-out occurs before a response is received, then a default response value (generally a default destination number) may be looked up in the customer record and put in an INAP/TCAP message and sent back to the requesting entity (steps 114 to 116). The SLP instance is then terminated (113).

Locating & Accessing Service Resources

The functionality associated with accessing a phone-page resource is schematically represented in FIG. 6 by resource access block 46. Block 46 includes URI determination block 47 for determining the URI of the phone page containing the desired resource on the basis of parameters passed to block 46. Using the URI returned by block 47, the resource access block 46 then accesses the phone page of the required service resource 49 over the Internet through interface 44.

Resource Codes—It is possible that more than one service resource is associated with a particular telephone number; in this case the resource access block 46 will need to know additional information (such as current point-in-call, pic) to enable the appropriate service resource to be identified. If the service resources associated with a number are located on different phone pages, then the additional information is also passed to the URI determination block 47 to enable it to return the URI of the appropriate phone page. It is also possible for all the service resources associated with a number to be located on the same phone page. In this case, the resource access block 46 uses the additional information to pass a resource-identifying parameter with its access request to the phone page concerned; it is then up to the functionality associated with the phone page to access the correct service resource.

Thus, each service resource can be considered as being identified by a respective resource code 54 (see FIG. 8) made up of a first part UI (“URI Identifier”) used to identify the URI at which the resource is located on the Internet, and a second part RRI (“Relative Resource Identifier”) used to identify the resource amongst plural resources at the same URI.

Resource Access—Where only one service resource 49 is located on a phone page 58 identified by a unique URI, then the resource code 54 simply comprises the UI, generally either a telephone number alone or a telephone number plus a pic parameter (see FIG. 9). In this case, accessing a resource simply involves mapping the whole resource code 54 into the corresponding URI (process 55) and then sending a request 57 to the corresponding phone page 58, this latter itself constituting the desired service resource 49. The result of accessing resource 49 is then returned in response message 59.

In contrast, where multiple service resources 49 are located on the same phone page 58 (FIG. 10), the resource code 54 comprises both a UI and RRI, the UI generally being a telephone number and the RRI a pic or other parameter for distinguishing between the co-located resources. In this case, accessing a resource involves mapping the UI part of the resource code 54 into the corresponding URI (process 55) and then sending a request 57 to the corresponding phone page (process 56), the request including the RRI of the resource code. The phone page 58 includes functionality 64 for accessing the required resource on the basis of the

US 7,012,916 B2

15

RRI in the request message. The result of accessing the required resource 49 is then returned in response message 59.

An alternative to the FIG. 10 method of accessing a service resource that is co-located with other resources on a phone page, would be to retrieve the whole page across the Internet simply using the URI derived from the UI part of the resource code, and then to extract the desired resource on the basis of the RRI.

URI Determination from Resource Code—The implementation of the URI determination block 47 that performs process 55 will next be considered. Block 47 may be implemented in a variety of ways, four of which are described below:

Direct Input

It would be possible, though not necessarily convenient, to arrange for the calling party to input directly the required URI. The calling party may thus input the host id component of the URI required (either in the form of a host domain name or host IP address) plus the path component of the URI. For example, in the case where the phone page of a called party is to be accessed, the calling party may input the URI of the called party and, indeed, this input may substitute for the normal input of a telephone number. A leading input string (for example “999”) may be used to identify the input as an URI. As regards the input means, where a user only has a standard 12 key telephone, input of host domain names and other URI elements requiring alpha characters, will need to be done using one of the standard techniques for alpha input from a phonepad (such techniques are already used, for example, to enable a calling party to “spell” out the name of the called party). It would also be possible to provide users with a full alphanumeric keypad to facilitate URI input.

Computation

Service resource access over the Internet could be restricted to a set of dialled numbers from which it was possible to compute a corresponding URI; in this case, this computation would be the responsibility of block 47.

Association Table Lookup

Probably the simplest implementation for the block 47 is as an association table (either in memory or held on database disc store 48) associating a URI with the UI part of each resource code. A potential problem with this approach is that a service resource may be required for a called party number on the other side of the world which implies a rigorous update regime between PSTN operators worldwide in order to keep the association table up-to-date. (Note that the same implication is not necessarily applicable in respect of marking the called-party number as one required to trigger a service request, since the number may be arranged to be one of a group of numbers all triggering an appropriate service request, in a manner similar to 800 numbers).

DNS-Type Lookup

An alternative lookup solution is to use a hierarchically-structured distributed database system, similar to (or even part of) the Domain Name System (DNS) of the Internet, in order to resolve the UI part of a resource code to a corresponding URI. This approach, which will be described in more detail below, would typically involve databases maintained by each PSTN operator for its numbers with which URIs are associated. These databases would be accessible by all PSTNs through a network such as the Internet with resolution requests

16

being pointed to the appropriate database in a manner similar to the Domain Name System. In this case, the block 47 is constituted by an appropriate resolution program arranged to request UI resolution over the Internet through interface 44.

Before describing a DNS-type lookup implementation for the URI determination block 47, some further general comments are appropriate. Whatever method is used to determine the URI, certain simplifications are possible if limited constraints are placed on the URIs permitted. In particular, it is not necessary to determine all components of an URI in the following cases:

- (i) A part of the URI path component can be made standard for all service resources, this standard part being simply added by the block 47 once the rest of the URI has been determined. For example, where a roaming number is to be looked up, it may by convention always be held in a file “roam” in a subdirectory “tel” of a subscriber’s directory on a particular server. In this case the URI host component and the subscriber-unique part of the path component are first determined and then the remaining path part “/tel/roam” is added.
- (ii) The URI path component can be arranged to be the same as a predetermined part of the resource code, the block 47 needing only to determine the host component and then add the path. For example, it may be agreed that the path must always end with the telephone number concerned, or sufficient of the terminating digits to have a high probability of uniqueness on the host machine. The path may also include standard components to be added by block 47.
- (iii) Blocks of telephone numbers may have their corresponding service resources located on the same host server so that it is only necessary to use a part of the telephone number to determine the host component of the URI; in this case, the path component can conveniently include all or part of each telephone number. This situation implies a tight degree of control by the telephone operators and does not offer the telephone user the freedom to choose the host server on which user places their phone page.

Another general point worthy of note is that however the URI is determined, the host component of the URI may be provided either in the form of a host domain name or a host IP address. Where the host is identified by a domain name, then a further resolution of URI host name to IP address will subsequently be carried out in standard manner by interface 44 using the Domain Name System of the Internet. This further resolution can be avoided if the host identity is directly provided as an IP address.

Where a database lookup is used to provide the number to URI translation, this database may be independent of, or combined with, a customer database containing other customer-related information. Factors affecting this choice include, on the one hand, the possible desirability of having the number-to-URI translation information widely available, and on the other hand, the possible desirability of restricting access to other customer-related information.

DNS-Type URI Lookup

A DNS-type lookup implementation for the URI determination block 47 will now be described in some detail for the case where the UI part of the resource code is a telephone number and there are no constraints on the URI, thereby requiring both the full host and path components of the URI to be returned by the lookup. A key part of the overall process is the formation of the equivalent of a host domain name from the telephone number of interest; this domain-

US 7,012,916 B2

17

name equivalent is then resolved into a corresponding URI by a lookup mechanism which in the present example is identical to that employed by the DNS (indeed, the lookup mechanism may be incorporated into the DNS though it can also be independently implemented).

The nature of the DNS has already been described above with reference to FIG. 3 when the term “DNS-type” system was also introduced. For convenience in the following a DNS-type system organised to provide a telephone number to URI translation facility will be referred to as a “Duris” system (standing for “DNS-type URI Server” system).

The basic principles surrounding operation of a Duris system are:

every telephone number can be turned into a host domain name (the name space containing such host domain names for the telephone numbers of interest is referred to below as the “telname space”); and

for every host domain name in the host domain space there is a Registration Record held by the Duris system containing the corresponding URI.

Thus, an input telephone number forming, in the present case, the UI part of a resource code 54 (see FIG. 11), is first parsed to form a domain name (step 120) and then passed to the Duris system (illustrated in FIG. 11 as formed by the DNS itself) to retrieve the RR with the corresponding URI (step 121). Following on from the URI lookup, if the URI returned has its host component as a domain name, the DNS is next used to derive the host IP address (step 122); this step is, of course not needed if the host component is stored as an IP address in the RR. The URI is then used to make a resource request to the appropriate server, passing any RRI part of the resource code 54 (step 123).

There are a number of possibilities at the top level as to how a Duris system could be implemented:

(a) Independent of the DNS. In this option, the telname space constitutes the entire name space to be managed with the root of the telname space being the “.” name space root (see FIG. 12A where the telname space is shown hatched). In this case, the Duris system is independent of the DNS itself. The Duris system could, of course, use the same basic infrastructure as the DNS (that is, the Internet) or an entirely separate network. Where the telname space comprises all the domain names corresponding to all public telephone numbers worldwide, parsing a full international telephone number would give a fully qualified domain name. Of course, the telname space could be a much smaller set of names such as those derived from internal extension numbers within a company having worldwide operations.

(b) Unfragmented Telname Space within the DNS. In this option, the telname space is a domain of the DNS name space and the Duris system is provided by the DNS itself. Thus, where the telname space comprises all domain names derived from public telephone numbers worldwide, the telname space could be placed within the domain of the ITU, in a special subdomain “tel”, the root of the telname space then being “tel.itu.int.” (see FIG. 12B where again, the hatched area represents the telname space). The responsibility for administering the domain “tel.itu.int.” would then lie with the ITU. With this latter example, to form a fully qualified domain name from an input telephone number, after the number has been parsed to form the part of the domain name corresponding to the structuring within the telname space, the tail “tel.itu.int.” is added. The fully qualified domain name is then applied to the DNS and

18

the corresponding RR record, holding the required URI, is retrieved. As a further example, the telname space could be all name derived from internal extension numbers within Hewlett-Packard in which case the root of the telname space would be “tel.hp.com.” and Hewlett-Packard would be entirely responsible for managing this domain.

(c) Fragmented Telname Space within the DNS. In this option, the telname space is split between multiple domains of the DNS name space and the Duris system is provided by the DNS itself. Thus where the telname space comprises all domain names derived from public telephone numbers worldwide, the telname space could be split between respective “tel” subdomains of each country domain; thus, as illustrated in FIG. 12C, the part of the telname space corresponding to French telephone numbers would have a root of “tel.fr.” and the part of the telname space corresponding to UK telephone numbers would subdomain would then lie with each country. With this latter example, to form a fully qualified domain name from an input telephone number, the part of the telephone number following the country code is parsed to form the part of the domain name within a country ‘tel’ subdomain and then a host domain name tail is added appropriate for the country concerned. Thus for a French telephone number input from X company will be parsed to a fully qualified domain name terminating “tel.yco.com.” and vice versa.

Consideration will next be given to the parsing of a telephone number into a domain name—in other words, where to insert the “.” characters into the number to provide the structuring of a domain name. Of course, as already explained, telephone numbers are hierarchically structured according to each country’s numbering plan. Thus one approach would be to follow this numbering plan structuring in dividing up a telephone number to form a domain name. By way of example, the telephone number “441447456987” which is a UK number (country code “44”) with a four digit area code (“1447”) and six digit local number (“456987”) could be divided to form a domain name of 456987.1447.44 (note that the reversal of label order occasioned by the fact that the DNS labels are arranged least significant first). If the telname space is a subdomain of the DNS with a placement as illustrated in FIG. 12B, the fully qualified domain name derived from the telephone number would be:

456987.1447.44.tel.itu.int.

There are however, difficulties inherent with trying to match the numbering plan hierarchy when parsing a telephone number into a host name. Firstly, in order to parse an international number correctly, it would be necessary for each entity tasked with this operation to know the structuring of each country’s numbering plan and where, as in the UK, area codes may be of differing length the required knowledge may need to take the form of a lookup table. Whilst this is not a complicated computational task, it is a major administrative nuisance as it means that each country will need to inform all others about its numbering plan and any updates. The second problem is that a six or seven digit local number is a very large domain; it would be preferable to create subdomains for performance and scaling reasons but there is no obvious way of doing this.

These problems can be overcome by giving up the restriction that the parsing of telephone number into a domain name should match the structuring of national numbering plans. In fact, there is no strong reason to follow such a scheme as DNS servers know nothing about the meaning of

US 7,012,916 B2

19

the name space. It is therefore possible to parse telephone numbers using a deterministic algorithm taking, for example, 4 digits at a time to limit the size of each subdomain and making it possible to 'insert the dots' without knowing the numbering plan concerned. So long as the DNS domains and zones served by the DNS servers are created correctly it will all work.

For international numbers it would still seem appropriate to separate off the country codes and so a hybrid parsing scheme would be to parse the initial part of a dialled number according to known country codes and thereafter use a deterministic scheme (for example 3,7 or 4,6 or 3,3,4) to separate the digits. Of course, if a fragmented telname space is being used as illustrated in Figure UC then the country code is used to look up the host name tail and it is only the national part of the number which would be parsed.

Finally, as regards the details of how a DNS server can be set up to hold RR records with URIs, reference can be made, for example, to "DNS and BIND", Paul Albitz and Cricket Liu, O'Reilly & Associates, 1992 which describes how to set up a DNS server using the Unix BIND implementation. The type of the RR records is, for example, text.

It should be noted that DNS labels should not in theory start with a digit. If this convention is retained, then it is of course a trivial exercise when parsing a telephone number to insert a standard character as the first character of each label. Thus, a 4 digit label of 2826 would become "t2826" where "t" is used as the standard starting character.

It will be appreciated that as with domain names, where an input telephone number is not the full number (for example, a local call does not require any international or area code prefix), it would be parsed into a domain name in the local domain.

The foregoing discussion of Duris system implementation, has been in terms of translating a telephone number into an URI where the telephone number forms the full UI of a resource code and the Duris system returns a full URI. It will be appreciated that the described Duris implementation can be readily adapted to accommodate the various modification discussed above regarding the form of the UI and what parts of the URI need to be looked up. For example, where there are a number of different service resources associated with a subscriber each in its own file and the required source is identified by a pic part of the resource code, then the input telephone number will be used to look up, not the full URI, but the host component and that part of the path component up to the relevant subdirectory, the pic part of the UI then being appended to identify the required resource file.

For small local Duris implementations, it may be possible to have a single server; the implementation should still, however, be considered as of a DNS type provided the other relevant features are present.

Nature of Service Resources

Turning now to a consideration of the service resources 49, how these service resources can be provisioned onto the servers 51 will be described more fully below but, by way of present example, the service resource or resources associated with a particular PSTN user (individual or organisation, whether a calling or called party) can be placed on a server 51 over the Internet from a user terminal 53 in one or more WWW pages.

Consider the simple case where the service resource is a service data item such as a telephone number (for example, an alternative number to be tried if the user's telephone corresponding to the number dialled by a calling party is busy). This diversion number could be made the sole service

20

resource of a phone page of the user. The phone page URI could be a URL with scheme set to HTTP in which case the GET method could be used to retrieve the diversion number. Such an arrangement is suitable if the phone page is only to be used for functional retrieval of the diversion number. However, if the diversion number is to be visually presented at a user terminal 53, then it may be desirable to accompany the number with explanatory material (this will often not be necessary as the diversion number can be arranged to be returned into an existing displayed page that already provides context information). However, where the phone page does include explanatory material as well as the diversion number, an entity only wishing to make functional use of the phone page, could be arranged to retrieve the phone page and then extract the diversion number (this would, of course, require a standard way of identifying the information to be extracted from the phone page).

An alternative and preferred arrangement for providing for both viewing and functional access to a resource requiring explanatory material for viewing, is to use an object-oriented approach to resource design. In this case, the resource object would have two different access methods associated with it, one for purely functional use of the resource and the other enabling viewing of associated explanatory material. It would then be up to the accessing entity to access the resource object using the appropriate object method.

Yet another arrangement for providing for both viewing and functional use of the diversion number, would be to provide separate resources appropriately configured for each use, each resource having its own resource code (generally, both such resources would be placed on the same phone page and in this case the UI part of each resource code would be the same).

Retrieval of a phone page for use by a human user will generally not be as time critical as retrieval for operational use by a PSTN. Thus, while for human use the scheme specified in the URL of a service resource could be HTTP, it may be advantageous for operational use to define a special "phone" scheme (access protocol) which would result in the server 51 using an optimised access routine to access the required resource (diversion number, in the current example) and respond to the accessing entity in the minimum possible time.

Besides data items, other possible types of service resource include service logic for execution in place (at the server) with the result of this execution being returned to the entity accessing the resource; service logic downloadable from the server to the accessing entity for execution at that entity; and a logging resource for logging information passed to it by the accessing entity (or simply for logging the fact that it has been accessed). It will be appreciated that the logging resource is really just a particular case of service logic executable in place.

By way of example, a service resource constituted by execute-in-place service logic can be arranged to implement time-of-day routing, the result of executing the service logic being the telephone number to which a call should be routed taking account of the time of day at the called party's location. An example of a service resource constituted by downloadable service logic is service logic for controlling calling-party option interrogation using the facilities provided by an IP. As regards the logging resource, this can be used for recording the number of calls placed to a particular number.

Where each resource has its own phone page and the resource is present only in its unembellished functional

US 7,012,916 B2

21

form, then the HTTP scheme can be employed for access using the GET method for both the downloadable service logic and the execution-in-place service logic, and the POST method for the logging resource. If it is desired to provide an explanatory material with each service resource, then any of the solutions discussed above in relation to data items, can be used.

Where more than one service resource is to be associated with a number, then each such resource can be placed on a respective phone page with its own URI. However, the preferred approach is to place all such service resources on the same page and use the RRI part of the corresponding resource codes to enable access to the appropriate resource. The accessed resource is then treated according to its form (executed if execute-in-place service logic, returned if downloadable service data or logic).

Thus if both a diversion-number service-data resource and a time-of-day execution-in-place service-logic resource are placed on the same phone page, the diversion-number resource code might have an RRI of "1" whilst the time-of-day resource code might have an RRI value of "2".

Where calling/called party options are to be included in a service resource for presentation to such party, then as already indicated, this can conveniently be done by constituting the service resource as downloadable service logic with the chosen option possibly initiating request for a follow-up service resource.

It will be appreciated that a service resource will often be of a complex type, combining service data and/or downloadable service logic and/or execute in place service logic. A particularly powerful combination is the combination of the two types of service logic where the downloadable service logic is designed to interact with execute-in-place service logic; using this arrangement, the user can be presented with complex client-server type applications.

Example Usage of Service Resource

FIG. 13 illustrates the operation of a service making use of a resource on a server 51. This service is equivalent to a "personal number" service by which a user can be accessed through a single, unchanging number even when moving between telephones having different real numbers. To achieve this, the user requiring this service (user B in the current example) is allotted a unique personal number (here referred to as the "Webtel" number of B) from a set of numbers all of which have the same leading number string to enable an SSP to readily identify a dialled number as a Webtel number. User B has a service resource 49 on a dedicated phone page on HTTP server 51, this phone page being located at a URL here identified as "URL (B phone page)". B's phone page when accessed returns the current roaming number ("B-telNb") where B can be reached. In the simplest case, B's phone page is just a single number that can be modified by B (for example, from a terminal 53) as B moves to a different phone. More likely is that B's phone page is an execute-in-place service logic providing time of day routing.

In the present example, the association between B's Webtel number and the URL of B's phone page is stored in an association table accessible to SCP 43.

Upon a user A seeking to contact user B by dialling the Webtel number of B, the telephone 40 being used by A passes a call set up request to SSP 41 (note that in FIG. 13 the bearer paths through the telephony network are shown by the thicker lines 60, the other heavy lines indicating signalling flows). SSP 41 detects the dialled number as a Webtel number and sends a service request to SCP 43 together with B's Webtel number. SCP 43 on receiving this

22

service request initiates a service logic program for controlling translation of B's Webtel number into a current roaming number for B; in fact, in the present case, this program simply requests the resource access block 46 to access the service resource identified by B's Webtel number, (that is, B's phone page 49) and return the result of this access. To this end, block 46 first translates B's Webtel number into the URL of B's phone page and then uses this URL to access B's phone page over the Internet (for example, using the 'phone' scheme already referred to with a method corresponding to the HTTP GET method). This results in B's current roaming number B-telNb being passed back to block 46 and in due course this number is returned to the SSP 41 which then initiates completion of call set up to the telephone 40 corresponding to B-telNb.

The FIG. 13 example related to a called-party service; it will, of course, be appreciated that the principle of accessing service resources over the Internet can be applied to all types of services, including both calling-party and called-party services and hybrids. Thus, standard 800 number services can be implemented with the dialled 800 number resulting in access to a phone page resource constituted by execute-in-place service logic that returns the most appropriate number for controlling onward call routing.

It will be appreciated that although in the FIG. 13 example the service request from the SSP was triggered by a leading number string of a dialled number, a service request may be triggered by a variety of triggers including calling-party number, called-party number, or some other user input, such triggers being possibly qualified by call setup progress (for example, called-party number qualified by a busy status or by ringing for more than a certain time).

With respect to the logging service resource mentioned above, one possible application for such a resource is in telephone voting. In this case, dialling the voting number causes the SSP picking up the call to pass a service request to SCP 43 which then contacts the appropriate logging resource over the Internet to register a vote after which the call is terminated. To minimise bottlenecks, a logging resource could be provided at a different URL for each SCP, it being a simple matter to collect and collate voting from all these logging resources over the Internet. If an SCP with Internet access is provided at every SSP, then the risk of congestion is greatly reduced.

As already noted, a user's phone page may hold multiple service resources in which case the access request from the accessing SCP needs to contain an appropriate RRI identifying the required resource.

In the event that an SCP is to provide both a traditional IN service to some users and an equivalent service using an Internet-accessed service resource to other users, then a lookup table may need to be provided in the SCP to ensure that a service request is appropriately handled; such a lookup table can conveniently be combined with the customer record database.

Once a user, such as user B, has set up one or more phone pages specifying his desired service resources (particularly service logic defining personalised services), it is clearly logical for user B to want any PSTN operator he cares to use, to access and utilise such service resources. This is possible if the Webtel-to-URI databases are available to all operators. Thus multiple operators could be set to access B's phone page or pages. If an operator declines to use B's phone pages, B can obviously choose not to use that operator (at least where that operator provides a long haul carrier service subject to user selection). The possibility therefore arises that service provision will cease to command a premium

US 7,012,916 B2

23

from operators, but that the provision of phone-page utilisation by an operator will become a necessary basic feature of PSTN operation.

Provisioning and Updating Service Resources

Consideration will next be given as to how the service resources 49 are provisioned to the servers 51 and subsequently updated.

So far as provisioning is concerned, two basic actions are required: firstly, the service resource must be placed on a server 51 and, secondly, the URI of the service resource must be notified to the PSTN operator along with the trigger conditions (number plus any other condition such as point in call) calling for access to the resource; if multiple resources are provided at the same URI, then the RRI values needed to retrieve the appropriate resource for a particular trigger condition, must also be notified. This notification process will be referred to hereinafter as 'registering' the service resource with the PSTN operator; registration is, of course, necessary to enable the association tables used by SCP 43 to be set up and for trigger conditions to be set in SSPs 43. For certain services, such as that described above with reference to FIG. 13, it is not the user that supplies the triggering number (the Webtel number in the FIG. 13 example); instead, the PSTN operator allocates an appropriate number to the user as part of the registration process.

As to the process of placing a service resource on a server 51, how this is carried out will depend on the attitude of the PSTN operator to the possible effects of such service resources on operation of the PSTN. Where the service resource simply returns a data item to an accessing entity, then an operator may not be too concerned about possible errors (accidental or deliberate) in implementing the service resource. However, the operator will probably be much more concerned about the proper operation of any service logic that may be returned by a resource; indeed, an operator may not permit such a service resource.

Assuming for the moment that an operator has no concerns about the nature or implementation of service resources, then how a resource is placed on a server 51 will largely depend on the nature of the server concerned. For example, if a user has a computer with network access to the Internet and this computer is used as server 51, then the user can simply load a desired resource onto the server as a WWW phone page for external access. A similar situation arises if the server is an organisation server to which the server has access over an internal LAN. In both these latter cases, loading the resource as a WWW phone page does not itself require Internet access. However, if the server 51 is one run by an external Internet service provider, then a user can arrange to download the required service resource into the user's allocated Web site space on the server; this may or may not involve Internet access. One special case of this latter scenario is where the PSTN operator provides a special server for user phone pages containing service resources.

Except where a user's own computer acts as server 51, placing a service resource on a server will generally involve clearing one or more levels of password protection. As regards the origin of the service resource loaded by a user onto server 51, this may be generated by the user or, particularly where the resource includes service logic, may be provided by a third party (including the PSTN operator).

If the PSTN operator wishes to have control over the service resources 49 to avoid any adverse effects on operation of the PSTN, two approaches are possible. Firstly, the operator could require that every resource (or, possibly, a particular subset) had to be subject to a verification process before use, appropriate measures then being taken to avoid

24

subsequent alteration of the resource by the user (except, possibly, for particular data items); in this respect, the operator could require that the resource be placed on a server under the operator's control and to which the user had no write access (except possibly for altering particular data items, as indicated above). A second, more attractive, approach to minimising adverse effects by the service resources 49, is for the operator to provide standard service resources to which a user could add the user's own data (and possibly make limited functional selections in case where the resource included service logic); the customised resource would then be loaded onto a server 51 controlled by the operator. This process can be conveniently implemented for a particular resource using an HTML "form" which a user could download over the WWW from the operator-controlled server. After completing the form and activating a 'commit' graphical button of the form, the entered information would be 'posted' back to the server where the information would be used to produce a customised service resource thereafter placed on the server for access over the Internet. An advantage of this approach is that registration of the service resource with the operator is simultaneously effected. (It may be noted that if registration needs to be done as a separate act from having a service resource loaded on a server, then using an HTML form is a very convenient way to implement the registration process).

From the foregoing it can be seen that whilst the provisioning process does not necessarily require information to be passed over the Internet, in many cases this will be the best solution, particularly if an HTML form exchanged over the WWW can be used to produce a customised service resource. It should be noted that producing a customised service resource using an HTML form is not limited to cases where the PSTN operator controls the server.

As regards updating service resources, there is likely to be a need to update certain data items on a fairly frequent basis (for example, roaming number). Where the PSTN operator does not place any controls on the service resources 49, then update is a relatively simple matter, only requiring write access to the server concerned (as already indicated, this will generally involve one or more levels of password protection). However, where the PSTN operator exercises control over the service resources, for example by only permitting customisations of standard service resources, such customised resources being loaded on servers controlled by the operator), then write access to the service resource may be tightly controlled. Again, an HTML form may conveniently be used as the medium for modifying a data item in such cases; to the operator, this has the benefit of limiting the modifications possible whilst to the user, a form interface should provide a simple route to resource modification.

For more complex updates, it may be necessary to go through a process similar to that required for initial provisioning.

Particularly where the service resources are held on a server 51 controlled by the PSTN operator, resource update will generally involve communication over the Internet.

Web User Interaction

Consideration will next be given to other possible uses of the service resources held in phone pages on the servers 51. For example, if user B's phone page contains a diversion number, then provided this phone page is read-accessible over the Internet from user A's terminal 53, user A can use a graphical Web browser running on terminal 53 to view B's phone page and discover B's diversion number. As earlier discussed, the diversion number may be passed to user A for

US 7,012,916 B2

25

display in an existing visual context giving meaning to the number, or may be passed to user A with accompanying explanatory text.

A more useful example is a current roaming number service for user B. Suppose B's phone page 49 on server 51 (see FIG. 14) is operative when accessed to return a current roaming number where B can be reached. Further suppose that user B has a Web site with several Web pages written in HTML and each page contains a graphical 'phone' button which when activated uses the GET method to access B's phone page by its URL. Now if user A whilst browsing (arrow 66) B's Web site over the WWW from user A's terminal 53, decides that he would like to call user B to discuss some item of interest, user A simply activates the phone button 65 on the currently viewed page of B. This causes B's phone page to be accessed using the HTTP request "GET URL (B Phone Page)"—see arrow 67.

B's current number to be called is then determined and passed to user A's terminal 53 (see arrow 68) where it is displayed. An explanatory text concerning the number will generally also be displayed; for example the text "Please call me at the following number:" could be displayed, this text being provided either by the HTML script associated with the phone button, or from the phone page when returning the current number. In fact, it would probably be more helpful to provide user A, not only with the current number for reaching user B, but also with all numbers where B could be reached together with the times when B was most likely to be at each number. Since this extra information is likely to be subject to frequent change, the only sensible way to provide the information is from the phone page. Thus, B's phone page not only provides the current number for reaching B, but also a text that includes numbers and times subject to change; scripting B's phone page is, of course, done in a way that ensures that variable data need only be altered in one place.

In a further example, B's phone page might include downloadable service logic for execution at user A's terminal. This is useful where choices are to be presented to a user, each choice producing a follow-up action such as fetching a further phone page. For example, the first-accessed phone page may be a family phone page giving the general telephone number for a family but also giving the user the possibility of selecting further phone information on each family member, such as a time-of-day dependent number; in this case, each family member has their own follow-up phone page.

In the above scenarios, user A has been presented with a number to call over the PSTN. User A can now pick up his standard telephone and dial the number given. In fact, a complication arises if A only has Internet access via a SLIP/PPP connection over an ordinary, non-ISDN, PSTN line since, in this case, A's telephone line is already tied up with making Internet access when gateway 90 seeks to set up a call to A's telephone; with an ISDN connection, as two channels are available, this problem does not arise. One way of overcoming this problem would be to have user A's terminal 53, after obtaining the number to call from B's phone page, automatically suspend its Internet session by storing any required state information (for example, current WWW URL being accessed) and then terminate its SLIP/PPP connection to thereby free up the telephone line. A can then telephone B. At the end of this call, A can resume the suspended Internet session, using the stored state information to return to the point where A left off to call B. An alternative approach is to operate a suitable multiplexing modulation scheme on the telephone line to A allowing

26

voice and data to be simultaneously carried. A number of such schemes already exist. The PSTN would then need to separate the combined data and voice streams coming from A at some point and pass each to its appropriate destination (the Internet data being forwarded to the ISP providing the SLIP/PPP connection for user A and the voice stream being passed to B); of course, data and voice traffic in the reverse direction would also need combining at some point for sending over the last leg to A's terminal.

Rather than A manually dialling B using a standard telephone, another possibility is that user A's terminal is provided with functionality enabling A to make a call over the PSTN from his terminal; this functionality generally comprises a hardware interface 70 (FIG. 14) to a telephone line and phone driver software 71 for driving the interface 70 in response to input from application software such as the Web browser 73. A could call up his phone software and enter the required number or, preferably, A need only "select" on screen the number returned from B's phone page and then pass it into A's phone software. Indeed, provided user B knew the software interface to the software 71 providing dialling functionality on A's terminal, it would be possible for B's phone page to return to A's terminal program code for automatically dialling B's number upon A confirming that he wishes to proceed with call placement. As an alternative to placing a voice call, if A's terminal is equipped with a suitable modem and controlling software, A could, instead, elect to send a fax or data to B through the PSTN either to B's ordinary number or to one specified in B's phone page as the number to be used for such transmissions. Of course, placing a call from A's terminal over the PSTN may be subject to the problem already discussed of conflict for use of the telephone line where this is not an ISDN line and A gains Internet access via a SLIP/PPP connection.

However the call is placed, if B's telephone corresponding to the number tried by A is busy, a number of possibilities exist. Thus if B has a phone page that specifies a diversion number, and B has registered this service resource with the PSTN, then the diversion number should be automatically tried by the PSTN. However, if the diversion number resource has not been registered with the PSTN, a busy signal will be returned to A. Where A has placed the call through a standard telephone, A must now decide how to proceed and A may elect either to give up or to refer again to B's phone page to look up the diversion number and redial using this number. If A placed the original call using his terminal 53 then the latter can be programmed to detect the return of a busy signal and then automatically look up B's diversion number and redial using this number. This functionality can be included in service logic downloaded from B's phone page and run on A's terminal.

If A had to terminate his Internet session in order to free up the telephone line for voice use, then referring back to B's phone page requires a new Internet session to be started (in fact, this inconvenience could be avoided if B's diversion number were passed to A's terminal at the time the original number to be dialled for B was supplied).

The service resource accessed on B's phone page upon B's telephone being busy may, of course, be more complex than just a diversion number. In particular, user A may be presented with a range of options including, for example, B's fax or voice mailbox number, the selection of an option potentially initiating the running of appropriate accessing software. Another possible option would be for A to leave B a call back message using a form downloaded from B's phone page upon this option being chosen; the completed

US 7,012,916 B2

27

form would be posted back to server **51** and logged for B to check in due course.

Of course, it may arise that user A wishes to access B's phone page to find out, for example, B's current roaming number, but user A does not know the URI of B's Web site and only has B's Webtel number. A could just call B through the PSTN in which case the translation of B's Webtel number to roaming number would be automatically effected (assuming B is still registered for this service); however, A may not wish to call B straight away, but just note his current roaming number. In order to solve A's problem, the Webtel-to-URI association tables previously described are preferably made accessible on the Internet at a known address (for example, at a known Web site). All that A need now do is to access this Web site passing B's Webtel number; B's phone page URI will then be returned to A who can then use it to access B's phone page. This process can, of course be made automatic from the point when A sends B's Webtel number to the association-table Web site.

Internet/PSTN Call Interface

In the FIG. **14** scenario, A's access to the PSTN was through a standard telephone interface even though the actual form of A's telephone differed from standard by being integrated into A's computer terminal **53**. FIG. **15** illustrates a situation where A, after being supplied with B's current roaming number as in the FIG. **14** case, calls B via a route that starts out over the Internet and then passes through a user network interface **80** into the PSTN. Interface **80** is arranged to convert between ISDN-type telephone signalling on the PSTN and corresponding signalling indications carried across the Internet in IP packets; in addition, interface **80** transfers voice data from IP packets onto trunk **60** and vice versa.

Thus, upon A initiating a call to B, Internet phone software **81** in A's terminal sends call initiation signalling over the Internet to interface **80**, the address of which is already known to A's terminal. At interface **80**, the signalling is converted into ISDN-type signalling and passed to SSP **41**. Call set up then proceeds in the normal way and return signalling is transferred back through interface **80**, over the Internet, to the software **81** in A's terminal. This software passes call setup progress information to the WWW browser **73** for display to A. Upon the call becoming established, A can talk to B through his telephone and A's voice input is first digitised in phone hardware interface **83** and then inserted into IP packets by software **81** to traverse the Internet to interface **80** (see arrow **84**); voice traffic from B follows the reverse path.

IN services can be provided to this call by SCP in response to a service request from an SSP **41**. Thus, if B's phone is busy, and B is registered for call diversion, SCP **43** on receiving a service request will access B's appropriate phone page for call diversion and retrieve the diversion number. If SSP **41** is not set to initiate a service request on B's telephone being busy, the busy indication is returned to A's terminal where it can be handled in the manner already described with reference to FIG. **14**.

In fact, interface **80** can be provided with functionality similar to an SSP to set trigger conditions and generate a service request to SCP **43** on these conditions being satisfied.

Third-Party Call Setup Gateway

FIG. **16** illustrates a further arrangement by which A can call B after receiving B's current roaming number. In this case, a third-party call set-up gateway **90** is provided that interfaces both with the Internet **50** and with an SSP **41**. Conveniently, gateway **90** can be co-located with SCP **43**

28

(though this is not essential). Gateway **90** has the capability of commanding SSP **41** to set up a call between specified telephones.

Thus, upon A wishing to call B, a third-party call setup request is sent from A's terminal over the Internet to the gateway **90** (see arrow **91**). This setup request includes A's telephone number and B's current roaming number. Gateway **90** first attempts to setup the call to A's telephone (which should generally succeed) and thereafter to set up the call to B's identified telephone. Once the call is setup, A and B communicate in standard manner across the PSTN.

If B's phone had been busy, then any of the previously described scenarios may ensue. Gateway **90** can also be arranged to make service requests to SCP **43** upon predetermined trigger conditions being satisfied. Thus, gateway **90** might be set to pick up the busy condition on B's telephone and initiate a service request to SCP **43** for a diversion number. However passing the busy indication back to A's terminal via gateway **90** is preferred because of the flexibility it gives A regarding further action.

As already generally discussed in relation to FIG. **14**, a complication arises if A only has Internet access via a SLIP/PPP connection over an ordinary, non-ISDN, PSTN line since, in this case, A's telephone line is already tied up with making Internet access when gateway **90** seeks to set up a call to A's telephone. The solutions discussed in respect of FIG. **14** (termination of Internet session; multiplexing voice and Internet data on same telephone line) can also be used here. An alternative approach both for FIG. **14** and for FIG. **16** scenarios is possible if user A's terminal can handle a voice call as digitised voice passed over the Internet. In this case, the voice call can be placed through an interface **80** of the FIG. **15** form, and the voice traffic and the Internet communication with the B's phone page and/or gateway **90** are both carried in Internet packets passed over the SLIP/PPP connection to/from A's terminal **53** but as logically distinct flows passed to separate applications running on terminal **53**.

It may be noted that the third party call setup request made by A's terminal to gateway **90** could equally have been made by service logic held in B's phone page and executed by server **51** (such an arrangement would, of course, require A's telephone number to be passed to B's phone-page service logic and this could be arranged to occur either automatically or through a form presented to user A at terminal A and then posted back to server **51**).

It may also be noted that the interface **80** of FIG. **15** and the gateway **90** of FIG. **16** provide examples of service requests being passed to the service control subsystem by entities other than SSPs **41**.

WWW-based "FreePhone" (800 Number) Services

It is possible to implement a "FreePhone" or "800 number" type of service using a combination of the WWW and the PSTN. As will be seen from the following description of such a service with reference to FIG. **17**, a WWW/PSTN implementation does not necessarily rely either on transferring call charges from the calling to called party or on the use of a special "800" number, two characteristics of standard "Freephone" schemes. The WWW/PSTN implementations do, however, possess the more general characteristic of placing an enquiring party and the party to whom the enquiry is directed, in telephone contact at the expense of the latter party.

In the FIG. **17** arrangement, a user D such as a large department store has a website on a server **51**; for the sake of simplicity, it will be assumed that the server is under the control of user D who has direct computer access to the

US 7,012,916 B2

29

server over line 125. D's Website may, for example, contain many catalogue-like Web pages illustrating goods offered for sale by D. In addition, D has a freephone page 124 for handling enquiries placed on a freephone basis; the URL of this page is associated with a "freephone" graphical button 122 placed on each of the Website catalogue pages.

Suppose user A at terminal 53 is browsing D's Website, looking at the catalogue pages (arrow 121). If A sees an item of interest and wishes to make an enquiry to D about this item, then A can activate at terminal 53 the graphical freephone button 122 associated with the catalogue page concerned. This activation causes code embedded in the catalogue page currently loaded in A's terminal to prompt the user to enter their telephone number and, optionally, their name, after which an HTTP request is sent to D's freephone page using the POST method and enclosing the entered data (arrow 123). D's freephone page on receiving this request executes service logic to enter a new enquiry (including A's name and telephone number) in an enquiry queue 127 maintained in an enquiry control system 126. In the present example, the enquiry control system is connected to the server 51 via line 125, externally of the Internet; however, it would also be possible to have server 51 communicate with the enquiry control system through the Internet and, indeed, this may be the most practical arrangement where D's Website is on an ISP server rather than on a server controlled by D. In fact, the code run in A's terminal upon activation of the freephone graphical button 122 could be arranged to directly forward the enquiry request to the enquiry control system over the Internet rather than passing it back through the server 51.

The enquiry control system 126 manages enquiries passed to it to ensure that they are dealt with in an ordered manner. The system 126 on receiving a new enquiry preferably estimates approximately how long it will be before the enquiry is dealt with, this estimation being based on the number of currently queued enquiries and the average time taken to handle an enquiry. This estimation of waiting time is passed back via server 51 to user A in the response to the POST request message.

The enquiry control system 126 looks after the distribution of enquiries to a number of agents each of which is equipped with a telephone 40 and a display 129. A's enquiry will be dealt with as soon as it reaches the head of the queue 127 and there is an agent detected as available to handle the enquiry (thus, for example, the system may be arranged to detect when an agent's telephone goes on hook). When these conditions are met, a distribution and setup control unit 128 takes A's enquiry and displays A's name and telephone number on the display 129 of the available agent (for clarity, herein referenced as agent D'); if user D keeps a database on D's past customers or credit rating data, then unit 128 will also look for and display any such further information known about A. At the same time, unit 128 makes a third-party call setup request (arrow 130) over the Internet to gateway 90 asking for a call to be set up between the telephone of the available agent D' and the telephone of user A, both telephones being identified by their respective numbers. If both D' and A pick up the call, the enquiry then proceeds, the cost of the call being paid for by D as it is D that originated the call over the PSTN. If, for whatever reason, the call remains incomplete (for example, unanswered by A) for a predetermined timeout period, then unit 128 can be arranged to automatically pass on to the next enquiry at the head of the queue 127.

It would, of course, be possible to dispense with having the unit 128 request call setup through gateway 90 and either

30

have the agent D' dial A's number manually or have unit 126 initiate auto-dialling for D' telephone (agent D' having, for example, a computer-integrated telephone similar to that of A's in FIG. 14). The advantage of these approaches is that the existing PSTN could be used without adaption and without any service installation, in implementing the WWW-based freephone service.

As discussed in relation to FIGS. 11 and 13, a complication arises in placing a call to A if A only has Internet access via a SLIP/PPP connection over an ordinary, non-ISDN, PSTN line since, in this case, A's telephone line is already tied up with making Internet access when user D tries to set up a call to A's telephone. The solutions discussed in respect of FIGS. 11 and 13 can also be used here (termination of Internet session; multiplexing voice and Internet data on same telephone line; and placing the call over the Internet to A's terminal). With respect to the solution based on termination of the Internet session, such termination could be delayed until A's enquiry was about to be dealt with; however, to do this, it would be necessary to provide feedback from the control system 126 over the Internet to A's terminal 53 and to associate this feedback with code for bringing about Internet-session termination. One way to achieve this would be to have the response message sent by server 51 in reply to the original POST request message from A, include a correlation code; any subsequent feedback from system 126 passed to A would also include this code (server A having also passed the code to control system 126) thereby allowing A's terminal to correctly identify this feedback. In fact, the same mechanism could be used to provide user A with updates on how much longer user A is likely to be waiting to be called back, this mechanism being usable independently of whether or not there was a conflict problem for use of A's telephone line.

Where user A only has a telephone 40 and no terminal 53, it is still possible to utilise the basic structure of FIG. 17 to provide a freephone service for user A without resorting to the complexity of call charge transfer. More particularly, A would dial a special number for user D's freephone service (typically an 800 number), and the SSP 41 would recognise this special number in standard manner and make a service request to SCP 43 including both this special number and A's number. SCP 43 would then ascertain D's freephone-page URL by doing a number-to-URL translation and access D's freephone page using a POST-method HTTP request similar to request 123. Once this request had been registered as an enquiry by D's freephone page 124, the latter could send a response to SCP 43 asking it to play an announcement such as "Your freephone enquiry has been registered; please hang up and you will be contacted shortly". This announcement could be played to A by an IP in standard manner. A would then hang up and be ready to receive a call from D.

A significant advantage of the above freephone schemes using WWW, is that user D is not running up charges for use of the PSTN during periods when an enquiry is enqueued, waiting to be handled.

Variants

Many variants are, of course, possible to the above-described arrangements and a number of these variants are described below.

Distributed Processing Environment. As is illustrated in FIG. 18, the SCP 43 may access the HTTP servers 51 through a distributed processing environment, DPE 98, at least logically separate from the Internet. Preferably in this case the servers 51 are controlled by PSTN operators and are thus restricted in number.

Service Resources on DNS-Type Servers. In the foregoing examples, the service resource items have been been

US 7,012,916 B2

31

placed on servers 51 connected to the Internet and a desired service resource has then been accessed over the Internet by the service control subsystem of the PSTN, and/or by Internet users, through the use of an URI derived from a resource code that identifies the the desired service resource item. In a preferred arrangement for deriving the URI from a resource code in the form of a telephone number, all or part of the telephone number concerned was parsed into domain name form and then resolved into an URI using a DNS-type distributed database system that, indeed, could be integrated into the DNS itself (see FIGS. 11 and 12, and related description). In fact, it would be possible to place service resource items directly in Registration Records held by a DNS-type distributed database system so that instead of the parsed telephone number being resolved to an URI which is then used to access the required resource, the parsed telephone number is directly resolved to the required service resource item. The mechanism employed in this process is exactly as already described for resolving a parsed telephone number into an URI. The DNS-type distributed database system used for this would preferably be one accessible over the Internet or the DNS itself so as to provide access to the service resource items for Internet users as well as for the service control subsystem of the PSTN (in the same manner as described above with reference to FIG. 18, the DNS-type servers holding the service resource items may be accessible to the service control subsystem by a network other than the Internet). Whilst the placing of service resource items in RRs held on DNS-type servers may not be suitable for all types of service resource items, it is suitable for items such as telephone numbers that do not change frequently. Thus, a suitable usage is to provide number portability; in this case, a dialled personal number triggers a lookup in the DNS-type system with all or part of the personal number being first parsed and then applied to the DNS type system to return a current number for call routing. All dialled numbers could be treated as personal numbers or simply a subset of such numbers, this subset comprising numbers that are readily identifiable as personal numbers by, for example, local lookup at an SSP or the presence of a predetermined leading digit string. The general concept of parsing a telephone number (or similar number) in whole or in part to form a domain name for resolution in a DNS-type distributed database system can be used for the retrieval of other items of information besides URIs and service resource items.

Feedback Mechanisms. In discussing the WWW-based freephone arrangement of FIG. 17, it was mentioned that user A could be supplied with feedback on the likely length of waiting time before A would be called back. This is one example of using the Internet to provide a feedback path for a potential or actual telephone user. Another example was provided in relation to FIG. 16 where the progress of call setup was reported back by the call setup gateway to user A's terminal. In fact, generally where a user is known to be using a terminal actively on the Internet the opportunity arises to provide the user with feedback on the progress of call setup through the telephone system. In order to do this, it is of course necessary to ensure that the feedback can be passed to the appropriate application running on terminal A and this will generally require the application to have made appropriate linking information available. As well as call setup progress information, other information can also be fed back for example during a call holding period. Thus, for example, a special server can be provided on the Internet holding multimedia clips or even videos that could be output to user A during a call holding period.

In the described arrangements, the servers 51 have held service resource items concerned primarily with call setup

32

control. It may be noted that in a somewhat different application, Internet servers could be arranged to hold data that could be accessed from the telephone system in response to a user-initiated telephone request and returned to that telephone user. Such a service would be provided, for example, in response to an SSP triggering a service request upon a particular telephone number being input, the service request prompting an SCP to cause an intelligent peripheral to access a particular Internet server (not necessarily an HTTP server) and retrieve the required data for return to the calling party. The intelligent peripheral may include a text-to-voice converter for replaying the data vocally to the user.

One further feedback process is also worthy of note, in this case in relation to service resource items themselves. By way of example, a telephone user G may subscribe to a service by which calls passed through to G's telephone are to be separated by a minimum of X minutes, X being user settable. To implement this service, G has a phone page on a server 51 that includes a "busy" status indication. Upon termination of a successful call to G, G's local SSP triggers the sending of a message by the associated SCP over the Internet to G's phone page. This message causes G's busy indication to be set to indicate that G is busy; the message also starts a timer which times out after a period X and causes the busy status indication to be reset. A call attempt to G will either be rejected at G's SSP because G's line is genuinely busy or will trigger the SSP to enquire via the SCP whether G's phone-page busy status indication is set. If the busy status indication is set (which it will be during the period X following termination of a successful call) the call attempt is rejected whereas if the busy status indication is in its reset condition, the call attempt is allowed to proceed. By placing the busy status indication mechanism on G's phone page, it is possible to arrange for G to be able to easily change the value of X.

More General Variants. Whilst the service control subsystem of the PSTN has been embodied as an SCP in the foregoing examples, it will be appreciated that the functionality of the service control subsystem could be provided as part of an SSP or in an associated adjunct. Furthermore, the triggering of service requests can be effected by equipment other than SSPs, for example by intercept boxes inserted in the SS7 signalling links.

It will be appreciated that the term "Internet" is to be understood to include not only the current specification of the TCP/IP protocols used for the Internet and the current addressing scheme, but also evolutions of these features such as may be needed to deal with isochronous media. Furthermore, references to the WWW and the HTTP protocol should equally be understood to encompass their evolved descendants.

The present invention can also be applied to telephone systems other than just PSTNs, for example to PLMNs and other mobile networks, and to private systems using PABXs. In this latter case, a LAN or campus-wide computer network serving generally the same internal users as the PABX, will take the role of the Internet in the described embodiments.

Furthermore, the present invention has application where any switched telecommunication system (for example, a broadband ATM system) requires service control and a computer network can be used for the delivery of service resources to the service control subsystem of the telecommunication system.

What is claimed is:

1. A method of accessing communication data relevant to a target entity identified by a number string, said method comprising:

US 7,012,916 B2

33

storing in the domain name system (DNS) of the Internet records each associated with a corresponding domain name and holding an URI for locating communications data associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a substantial portion of the number string into at least a part of said domain name; applying said process to the said number string identifying the target entity whereby to form the related domain name;

supplying the domain name formed to the DNS to retrieve the URI held in the corresponding record; and using the URI retrieved to access said communication data.

2. A method according to claim 1, wherein the URI held in at least one said record is an URL including an access scheme and host address.

3. A method according to claim 2, wherein said host address is a domain name.

4. A method according to claim 1, wherein each said number string is in telephone-number form.

5. A method according to claim 1, wherein the communications data is a communications endpoint address for the target entity.

6. A method of accessing a target entity over a telephone network, in which communications data in the form of a telephone number for the target entity is accessed according to the method of claim 1, this telephone number then being used to call the target entity over the telephone network.

7. A method according to claim 1, wherein the URI held in at least one said records is the URI for said communications data itself.

8. A method according to claim 1, wherein the URI held in at least one said records is of functionality that has access to multiple items of communications data, and wherein using the URI retrieved to access said communication data involves using the URI to access said functionality and supply it with an indicator of the desired item of communications data, this data then being returned by the functionality.

9. A method according to claim 8, wherein the indicator is incorporated into said URI and supplied in this form to said functionality.

10. A method according to claim 8, wherein the indicator is supplied to said functionality as a separate element to said URI.

11. A method according to claim 8, wherein the URI held in at least one said records is an URL including an access scheme and host address for accessing said functionality.

12. A method according to claim 8, wherein each said number string is in telephone-number form.

13. A method according to claim 8, wherein the communications data is a communications endpoint address for the target entity.

14. A method of accessing a target entity over a telephone network, in which communications data in the form of a telephone number for the target entity is accessed according to the method of claim 7, this telephone number then being used to call the target entity over the telephone network.

15. A method of accessing communication data relevant to a target entity identified by a number string, said method comprising:

storing in the domain name system (DNS) of the Internet records each associated with a corresponding domain name and holding an at least part-formed URL, including access scheme and host name, of an item of

34

communications data, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a substantial portion of the number string into at least a part of said domain name;

applying said process to the said number string identifying the target entity whereby to form the related domain name;

supplying the domain name formed to the DNS to retrieve the at least part-formed URL held in the corresponding record; and

using the at least part-formed URL retrieved to access said communication data.

16. A method according to claim 15, wherein the URL retrieved is only part-formed and using the at least part-formed URL retrieved to access said communication data involves completing the URL by adding a path element serving to distinguish the desired item of communication data from other items held on the same host.

17. A method according to claim 15, wherein each said number string is in telephone-number form.

18. A method according to claim 15, wherein the communications data is a communications endpoint address for the target entity.

19. A method of accessing a target entity over a telephone network, in which communications data in the form of a telephone number for the target entity is accessed according to the method of claim 15, this telephone number then being used to call the target entity over the telephone network.

20. A method of discovering communications endpoint address data for contacting a target entity identified by a number string, said method comprising:

storing in the domain name system (DNS) of the Internet records each associated with a corresponding domain name and holding an URL of a resource that has access to multiple items of communications endpoint address data, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a substantial portion of the number string into at least a part of said domain name, the number strings being in telephone-number form;

applying said process to the said number string identifying the target entity whereby to form the related domain name;

supplying the domain name formed in step (b) to the DNS to retrieve the URL held in the corresponding record; and

using the URL to access the corresponding said resource and supply it with an indicator of the desired item of communications endpoint address data, this data then being returned by the resource.

21. A method of accessing communication data relevant to a target entity identified by a number string, said method comprising:

storing in a DNS-type database system, records each associated with a corresponding domain name and holding an URI for locating communications data associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a substantial portion of the number string into at least a part of said domain name;

applying said process to the said number string identifying the target entity whereby to form the related domain name;

US 7,012,916 B2

35

supplying the domain name formed to the DNS-type database system to retrieve the URI held in the corresponding record; and
using the URI retrieved to access said communication data.

22. A method according to claim 21, wherein the URI held in at least one said records is an URL including an access scheme and host address for accessing said communications data.

23. A method according to claim 21, wherein each said number string is in telephone-number form.

24. A method according to claim 21, wherein the communications data is a communications endpoint address for the target entity.

25. A method according to claim 21, wherein the URI held in at least one said record is of functionality that has access to multiple items of communications data, and wherein using the URI retrieved to access said communication data involves using the URI to access said functionality and supply it with an indicator of the desired item of communications data which is then returned by the functionality.

26. A method of accessing communications data for contacting a target entity, said method comprising:

forming, from a number string identifying the target entity, a domain name by a process including parsing at least a substantial portion of the number string into at least a part of said domain name;

supplying the domain name formed to the domain name system of the Internet and receiving back from the domain name system a resource record including an URI for locating communications data associated with the domain name; and

using the URI received back to access said communication data.

27. A method according to claim 26, wherein the URI received back is an URL including an access scheme and host address for accessing said communications data.

28. A method according to claim 26, wherein said number string is a telephone number.

29. A method according to claim 26, wherein the communications data is a communications endpoint address for the target entity.

30. A method of accessing a target entity over a telephone network, in which communications data in the form of a telephone number for the target entity is accessed according to the method of claim 26, this telephone number then being used to call the target entity over the telephone network.

31. A method according to claim 26, wherein the URI received back is the URI of said communications data itself.

32. A method according to claim 26, wherein the URI received back is of functionality that has access to multiple items of communications data, and wherein using the URI received back to access said communication data involves using the URI to access said functionality and supply it with an indicator of the desired item of communications data which is then returned by the functionality.

33. A method according to claim 32, wherein the indicator is incorporated into said URI and supplied in this form to said functionality.

34. A method according to claim 32, wherein the indicator is supplied to said functionality as a separate element to said URI.

35. A method according to claim 32, wherein the URI received back is an URL including an access scheme and host address for accessing said functionality.

36. A method according to claim 32, wherein said number string is a telephone number.

36

37. A method according to claim 32, wherein the communications data is a communications endpoint address for the target entity.

38. A method of accessing a target entity over a telephone network, in which communications data in the form of a telephone number for the target entity is accessed according to the method of claim 32, this telephone number then being used to call the target entity over the telephone network.

39. A method of accessing communications data for contacting a target entity, said method comprising:

forming, from a number string identifying the target entity, a domain name by a process including parsing at least a substantial portion of the number string into at least a part of said domain name;

supplying the domain name formed to the domain name system of the Internet and receiving back from the domain name system a resource record including an at least part-formed URL, including access scheme and host name, of an item of communications data associated with the domain name; and

using the URL received back to access said communication data.

40. A method according to claim 39, wherein the URL received back is only part-formed and wherein using the URI received back to access said communications data involves completing the URL by adding a path element serving to distinguish the desired item of communication data from other items held on the same host.

41. A method according to claim 39, wherein said number string is a telephone number.

42. A method according to claim 39, wherein the communications data is a communications endpoint address for the target entity.

43. A method of accessing a target entity over a telephone network, in which communications data in the form of a telephone number for the target entity is accessed according to the method of claim 39, this telephone number then being used to call the target entity over the telephone network.

44. A method of discovering communications endpoint address data for contacting a target entity, said method comprising:

forming, from a number string identifying the target entity, a domain name by a process including parsing at least a substantial portion of the number string into at least a part of said domain name;

supplying the domain name formed to the domain name system of the Internet and receiving back from the domain name system a resource record including an URI of a resource that has access to multiple items of communications endpoint address data; and

using the URI received back to access corresponding said resource and supply it with an indicator of the desired item of communications endpoint address data, this data then being returned by the functionality.

45. A method of accessing communications data for contacting a target entity, said method comprising:

forming, from a number string identifying the target entity, a domain name by a process including parsing at least a substantial portion of the number string into at least a part of said domain name;

supplying the domain name formed to a DNS-type database system and receiving back a resource record including an URI for locating communications data associated with the domain name; and

using the URI received back to access said communications data.

37

46. A method according to claim 45, wherein the URI received back is an URL including an access scheme and host address for accessing said communications data.

47. A method according to claim 45, wherein each said number string is a telephone number.

48. A method according to claim 45, wherein the communications data is a communications endpoint address for the target entity.

49. A method according to claim 45, wherein the URI held in at least one said record is of functionality that has access to multiple items of communications data, and wherein using the URI received back to access said communications data involves using the URI to access said functionality and supply it with an indicator of the desired item of communications data which is then returned by the functionality.

50. A server of the domain name system of the Internet, the server holding at least one resource record that provides a mapping from a domain name to an URI for locating communications data associated with the domain name, at least a substantial portion of the domain name being in the form of a number string that has been parsed into plural domain-name labels.

51. A server according to claim 50, wherein said URI is an URL including an access scheme and host address for accessing said communications data.

52. A method according to claim 50, wherein the communications data is a communications endpoint address for the target entity.

53. A server according to claim 51, wherein said number string is at least a substantial portion of a telephone number.

54. A server according to claim 50, wherein said URI is of functionality that has access to multiple items of communications data.

55. A server of the domain name system of the Internet, the server holding at least one resource record that provides

38

a mapping from a domain name to an at least part-formed URL, including access scheme and host name, of an item of communications data associated with the domain name, at least a substantial portion of the domain name being in the form of a number string that has been parsed into plural domain-name labels.

56. A server according to claim 55, wherein said number string is at least a substantial portion of a telephone number.

57. A server of the domain name system of the Internet, the server holding at least one resource record that provides a mapping from a domain name to an URL of a resource that has access to multiple items of communications data, at least a substantial portion of the domain name being in the form of a number string that has been parsed into plural domain-name labels.

58. A server according to claim 57, wherein said number string is at least a substantial portion of a telephone number.

59. A DNS-type distributed database system holding at least one resource record that provides a mapping from a domain name to an URI for locating communications data associated with the domain name, at least a substantial portion of the domain name being in the form of a number string that has been parsed into plural domain-name labels.

60. A system according to claim 59, wherein said URI is an URL including an access scheme and host address for accessing said communications data.

61. A system according to claim 59, wherein said number string is at least a substantial portion of a telephone number.

62. A system according to claim 59, wherein said URI is of functionality that has access to multiple items of communications data.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 7,012,916 B2
APPLICATION NO. : 10/052285
DATED : March 14, 2006
INVENTOR(S) : Colin Low et al.

Page 1 of 1

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In the Title Page, item (75), the following named inventor should be removed:

Nicolas Bouthors

In the Title Page, item (30), the following foreign applications should be removed:

GB 9525190 filed December 11, 1995
EP 95410148 filed December 22, 1995

Signed and Sealed this

Twelfth Day of August, 2008

A handwritten signature in black ink, appearing to read "Jon W. Dudas". The signature is stylized with a large, looped initial "J" and a distinct "D".

JON W. DUDAS
Director of the United States Patent and Trademark Office

C.A. 12-0205 RGA
Comcast IP Holdings, LLC
v.
Sprint Commc'ns Co., et al.

**PTX
002**

U 7405603



THE UNITED STATES OF AMERICA

TO ALL TO WHOM THESE PRESENTS SHALL COME:

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

March 08, 2013


**THIS IS TO CERTIFY THAT ANNEXED HERETO IS A TRUE COPY FROM
THE RECORDS OF THIS OFFICE OF:**

U.S. PATENT: 8,170,008

ISSUE DATE: May 01, 2012

By Authority of the
Under Secretary of Commerce for Intellectual Property
and Director of the United States Patent and Trademark Office




GLORIA A MURRAY
Certifying Officer

COM-DEL(0205)-00141724

A108

PTX002-0001



US008170008B2

(12) **United States Patent**
Low et al.

(10) **Patent No.:** **US 8,170,008 B2**
(45) **Date of Patent:** **May 1, 2012**

(54) **METHOD AND APPARATUS FOR ACCESSING COMMUNICATION DATA RELEVANT TO A TARGET ENTITY IDENTIFIED BY A NUMBER STRING**

(52) **U.S. Cl.** 370/352; 709/229
(58) **Field of Classification Search** 370/352
See application file for complete search history.

(75) **Inventors:** **Colin Low, Gloucestershire (GB); Andrew Franklin Seaborne, Bristol (GB); Nicolas Bouthors, Les Bealires (FR)**

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,054,756 A 10/1977 Comella et al.
4,100,377 A 7/1978 Flanagan
4,201,891 A 5/1980 Lawrence et al.
4,313,035 A 1/1982 Jordan et al.

(Continued)

(73) **Assignee:** **Comcast IP Holdings I, LLC, Wilmington, DE (US)**

(*) **Notice:** Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 259 days.

FOREIGN PATENT DOCUMENTS

EP 0559979 9/1993

(Continued)

(21) **Appl. No.:** **12/389,843**

(22) **Filed:** **Feb. 20, 2009**

OTHER PUBLICATIONS

Berners-Lee, et al., "Hypertext Transfer Protocol—HTTP/1.0", HTTP Working Group Internet Draft (Oct. 1995).

(Continued)

(65) **Prior Publication Data**

US 2010/0046735 A1 Feb. 25, 2010

Related U.S. Application Data

(60) Continuation of application No. 11/066,880, filed on Feb. 25, 2005, now Pat. No. 7,804,816, which is a division of application No. 10/052,285, filed on Jan. 18, 2002, now Pat. No. 7,012,916, which is a division of application No. 09/077,795, filed as application No. PCT/GB96/03055 on Dec. 11, 1996, now Pat. No. 6,466,570.

Primary Examiner — Andrew Chriss

(74) *Attorney, Agent, or Firm* — Banner & Witcoff, Ltd.

(57) **ABSTRACT**

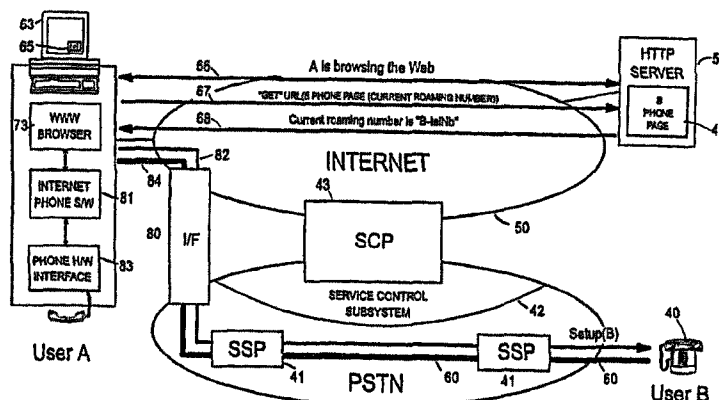
Service resource items for use in call setup in a telephone system are held on servers that are connected to a computer network which is logically distinct from the telephone system infrastructure; this computer network may, for example, make use of the Internet. Each service item is locatable on the network at a corresponding URI and is associated with a particular telephone number. A mapping is provided between telephone numbers and the URIs of associated service resource items. When it is desired to access a service resource item associated with a particular telephone number, this mapping is used to retrieve the corresponding URI which is then used to access the desired service resource item.

(30) **Foreign Application Priority Data**

Dec. 11, 1995 (GB) 9525190.6
Dec. 22, 1995 (EP) 95410148
Feb. 20, 1996 (GB) 9603582.9

(51) **Int. Cl.**
H04L 12/66 (2006.01)
G06F 15/16 (2006.01)

29 Claims, 14 Drawing Sheets



US 8,170,008 B2

Page 2

U.S. PATENT DOCUMENTS

4,611,094 A	9/1986	Asmuth et al.	5,452,289 A	9/1995	Sharma et al.
4,611,096 A	9/1986	Asmuth et al.	5,452,350 A	9/1995	Reynolds et al.
4,653,045 A	3/1987	Stanley et al.	5,455,821 A	10/1995	Schaeffer et al.
4,659,877 A	4/1987	Dorsey et al.	5,459,775 A	10/1995	Isono et al.
4,679,190 A	7/1987	Dias et al.	5,461,611 A	10/1995	Drake, Jr. et al.
4,685,125 A	8/1987	Zave	5,463,677 A	10/1995	Bash et al.
4,713,806 A	12/1987	Oberlander et al.	5,469,496 A	11/1995	Emery et al.
4,718,005 A	1/1988	Feigenbaum et al.	5,469,497 A	11/1995	Pierce et al.
4,734,931 A	3/1988	Bourg et al.	5,473,608 A	12/1995	Gagne et al.
4,747,130 A	5/1988	Ho	5,475,748 A	12/1995	Jones
4,771,425 A	9/1988	Baran et al.	5,475,817 A	12/1995	Waldo et al.
4,788,718 A	11/1988	McNabb et al.	5,477,531 A	12/1995	McKee et al.
4,821,034 A	4/1989	Anderson et al.	5,479,494 A	12/1995	Clithrow
4,872,157 A	10/1989	Hemmady et al.	5,481,603 A	1/1996	Gutierrez et al.
4,893,302 A	1/1990	Hemmady et al.	5,483,586 A	1/1996	Sussman
4,897,874 A	1/1990	Lidinsky et al.	5,483,587 A	1/1996	Hogan et al.
4,899,373 A	2/1990	Lee et al.	5,487,111 A	1/1996	Slusky
4,933,931 A	6/1990	Kokuho	5,495,521 A	2/1996	Rangachar
4,969,184 A	11/1990	Gordon et al.	5,500,889 A	3/1996	Baker et al.
4,979,206 A	12/1990	Padden et al.	5,500,929 A	3/1996	Dickinson
4,996,707 A	2/1991	O'Malley et al.	5,506,887 A	4/1996	Emery et al.
5,025,254 A	6/1991	Hess	5,506,893 A	4/1996	Buscher et al.
5,029,196 A	7/1991	Morganstein	5,511,111 A	4/1996	Serbetcioğlu et al.
5,065,393 A	11/1991	Sibbitt et al.	5,513,127 A	4/1996	Gard et al.
5,068,888 A	11/1991	Scherk et al.	5,521,719 A	5/1996	Yamada
5,107,492 A	4/1992	Roux et al.	5,521,924 A	5/1996	Kakuma et al.
5,113,499 A	5/1992	Ankney et al.	5,524,137 A	6/1996	Rhee
5,115,495 A	5/1992	Tsuchiya et al.	5,526,353 A	6/1996	Henley et al.
5,146,488 A	9/1992	Okada et al.	5,526,416 A	6/1996	Dezono et al.
5,146,491 A	9/1992	Silver et al.	5,526,489 A	6/1996	Nilakantan et al.
5,157,390 A	10/1992	Yoshie et al.	5,530,744 A	6/1996	Charalambous et al.
5,157,662 A	10/1992	Tadamura et al.	5,530,852 A	6/1996	Meske, Jr. et al.
5,159,592 A	10/1992	Perkins	5,533,115 A *	7/1996	Hollenbach et al. 379/221.11
5,159,624 A	10/1992	Makita	5,539,884 A	7/1996	Robrock, II
5,164,938 A	11/1992	Jurkevich et al.	5,539,886 A	7/1996	Aldred et al.
5,185,860 A	2/1993	Wu	5,541,917 A	7/1996	Farris
5,193,110 A	3/1993	Jones et al.	5,541,927 A	7/1996	Kristol et al.
5,195,085 A	3/1993	Bertsch et al.	5,541,930 A	7/1996	Klingman
5,195,086 A	3/1993	Baumgartner et al.	5,546,452 A	8/1996	Andrews et al.
5,204,894 A	4/1993	Darden	5,551,025 A	8/1996	O'Reilly et al.
5,206,901 A	4/1993	Harlow et al.	5,551,035 A	8/1996	Arnold et al.
5,223,699 A	6/1993	Flynn et al.	5,559,721 A	9/1996	Ishii
5,226,075 A	7/1993	Funk et al.	5,559,871 A	9/1996	Smith
5,241,588 A	8/1993	Babson, III et al.	5,561,670 A	10/1996	Hoffert et al.
5,247,571 A	9/1993	Kay et al.	5,563,882 A	10/1996	Bruno et al.
5,260,986 A	11/1993	Pershan	5,570,410 A	10/1996	Hooshiani
5,272,749 A	12/1993	Masek	5,575,961 A	11/1996	Kuwabara et al.
5,287,199 A	2/1994	Zoccolillo	5,579,472 A	11/1996	Keyworth, II et al.
5,311,583 A	5/1994	Friedes et al.	5,583,920 A	12/1996	Wheeler, Jr.
5,317,566 A	5/1994	Joshi	5,586,177 A	12/1996	Farris et al.
5,325,421 A	6/1994	Hou et al.	5,590,127 A	12/1996	Bales et al.
5,327,486 A	7/1994	Wolff et al.	5,590,181 A	12/1996	Hogan et al.
5,329,520 A	7/1994	Richardson	5,594,769 A	1/1997	Pellegrino et al.
5,329,578 A	7/1994	Brennan et al.	5,594,789 A	1/1997	Seasholtz et al.
5,333,266 A	7/1994	Boaz et al.	5,603,084 A *	2/1997	Henry et al. 455/419
5,347,633 A	9/1994	Ashfield et al.	5,603,085 A	2/1997	Shedlo
5,351,276 A	9/1994	Doll, Jr. et al.	5,604,682 A	2/1997	McLaughlin et al.
5,351,286 A	9/1994	Nici	5,604,737 A	2/1997	Iwami et al.
5,353,331 A	10/1994	Emery et al.	5,608,446 A	3/1997	Carr et al.
5,361,256 A	11/1994	Doeringer et al.	5,608,786 A	3/1997	Gordon
5,390,335 A	2/1995	Stephan et al.	5,610,910 A	3/1997	Focsaneanu et al.
5,392,344 A	2/1995	Ash et al.	5,610,972 A	3/1997	Emery et al.
5,396,542 A	3/1995	Alger et al.	5,610,976 A	3/1997	Uota et al.
5,402,478 A	3/1995	Hluchyj et al.	5,617,422 A	4/1997	Litzenberger et al.
5,406,557 A	4/1995	Baudoin	5,617,540 A	4/1997	Civanlar et al.
5,418,844 A	5/1995	Morrissey et al.	5,619,555 A	4/1997	Fenton et al.
5,422,940 A	6/1995	Endo et al.	5,619,557 A	4/1997	Van Berkum
5,423,003 A	6/1995	Berteau	5,619,562 A	4/1997	Maurer et al.
5,425,028 A	6/1995	Britton et al.	5,623,601 A	4/1997	Vu
5,425,085 A	6/1995	Weinberger et al.	5,625,404 A	4/1997	Grady et al.
5,425,091 A	6/1995	Josephs	5,625,407 A	4/1997	Biggs et al.
5,428,608 A	6/1995	Freeman et al.	5,625,675 A	4/1997	Katsumaru et al.
5,436,963 A	7/1995	Fitzpatrick et al.	5,625,677 A	4/1997	Feiertag et al.
5,438,568 A	8/1995	Weisser, Jr.	5,625,682 A	4/1997	Gray et al.
5,440,563 A	8/1995	Isidoro et al.	5,627,886 A	5/1997	Bowman
5,440,620 A	8/1995	Slusky	5,629,978 A *	5/1997	Blumhardt et al. 370/271
5,448,633 A	9/1995	Jamaledin et al.	5,633,916 A	5/1997	Goldhagen et al.
5,450,411 A	9/1995	Heil	5,634,012 A	5/1997	Stefik et al.
			5,636,216 A	6/1997	Fox et al.

US 8,170,008 B2

Page 3

5,646,982 A	7/1997	Hogan et al.	5,825,772 A	10/1998	Dobbins et al.
5,651,006 A	7/1997	Fujino et al.	5,825,869 A	10/1998	Brooks et al.
5,652,787 A	7/1997	O'Kelly	5,826,268 A	10/1998	Schaefer et al.
5,654,957 A	8/1997	Koyama	5,828,737 A	10/1998	Sawyer
5,657,250 A	8/1997	Park et al.	5,828,844 A	10/1998	Civanlar et al.
5,661,197 A	8/1997	Villiger et al.	5,832,197 A	11/1998	Houji
5,661,782 A	8/1997	Bartholomew et al.	5,835,723 A	11/1998	Andrews et al.
5,661,790 A	8/1997	Hsu	5,838,682 A	11/1998	Dekelbaum et al.
5,664,005 A	9/1997	Emery et al.	5,838,686 A	11/1998	Ozkan
5,664,102 A	9/1997	Faynberg	5,838,768 A	11/1998	Sumar et al.
5,668,857 A	9/1997	McHale	5,838,970 A	11/1998	Thomas
5,673,263 A	9/1997	Basso et al.	5,845,267 A	12/1998	Ronen
5,675,507 A	10/1997	Bobo, II	5,848,143 A	12/1998	Andrews et al.
5,675,741 A	10/1997	Aggarwal et al.	5,850,433 A	12/1998	Rondeau
5,680,116 A	10/1997	Hashimoto et al.	5,864,604 A	1/1999	Moen et al.
5,680,392 A	10/1997	Semaan	5,864,610 A	1/1999	Ronen
5,687,167 A	11/1997	Bertin et al.	5,864,670 A	1/1999	Hayashi et al.
5,689,550 A	11/1997	Garson et al.	5,867,495 A	2/1999	Elliot et al.
5,689,553 A	11/1997	Ahuja et al.	5,867,566 A	2/1999	Hogan et al.
5,692,039 A	11/1997	Brankley et al.	5,867,571 A	2/1999	Borchering
5,694,463 A	12/1997	Christie et al.	5,870,454 A	2/1999	Dahlen
5,699,089 A	12/1997	Murray	5,870,561 A	2/1999	Jarvis et al.
5,699,352 A	12/1997	Kriete et al.	5,873,077 A	2/1999	Kanoh et al.
5,701,295 A	12/1997	Bales et al.	5,878,126 A	3/1999	Velamuri et al.
5,701,465 A	12/1997	Baughner et al.	5,878,128 A	3/1999	Kantola
5,703,935 A	12/1997	Raissy et al.	5,878,212 A	3/1999	Civanlar et al.
5,703,940 A	12/1997	Sattar et al.	5,881,134 A	3/1999	Foster et al.
5,703,942 A	12/1997	Pinard et al.	5,883,891 A	3/1999	Williams et al.
5,708,780 A	1/1998	Levergood et al.	5,883,894 A	3/1999	Patel et al.
5,710,884 A	1/1998	Dedrick	5,884,032 A	3/1999	Bateman et al.
5,712,903 A	1/1998	Bartholomew et al.	5,889,774 A	3/1999	Mirashrafi et al.
5,712,906 A	1/1998	Grady et al.	5,889,777 A	3/1999	Miyao et al.
5,712,907 A	1/1998	Wegner et al.	5,892,754 A	4/1999	Kompella et al.
5,712,908 A	1/1998	Brinkman et al.	5,903,558 A	5/1999	Jones et al.
5,719,854 A	2/1998	Choudhury et al.	5,903,559 A	5/1999	Acharya et al.
5,724,355 A	3/1998	Bruno et al.	5,903,735 A	5/1999	Kidder et al.
5,724,406 A	3/1998	Juster	5,907,547 A	5/1999	Foladare et al.
5,724,412 A	3/1998	Srinivasan	5,909,430 A	6/1999	Reaves
5,726,984 A	3/1998	Kubler et al.	5,915,008 A	6/1999	Dulman
5,727,129 A	3/1998	Barrett et al.	5,915,012 A	6/1999	Miloslavsky
5,729,544 A	3/1998	Lev et al.	5,917,817 A	6/1999	Dunn et al.
5,729,599 A	3/1998	Plomondon et al.	5,923,659 A	7/1999	Curry et al.
5,732,078 A	3/1998	Arango	5,930,348 A	7/1999	Regnier et al.
5,732,219 A	3/1998	Blumer et al.	5,933,412 A	8/1999	Choudhury et al.
5,737,333 A	4/1998	Civanlar et al.	5,933,425 A	8/1999	Iwata
5,737,395 A	4/1998	Iribarren	5,933,486 A	8/1999	Norby et al.
5,737,414 A	4/1998	Walker et al.	5,936,958 A	8/1999	Soumiya et al.
5,742,668 A	4/1998	Pepe et al.	5,940,479 A	8/1999	Guy et al.
5,742,905 A	4/1998	Pepe et al.	5,940,616 A	8/1999	Wang
5,744,533 A	4/1998	Iwamoto et al.	5,944,795 A	8/1999	Civanlar
5,745,694 A	4/1998	Egawa et al.	5,949,869 A	9/1999	Sink et al.
5,748,468 A	5/1998	Notenboom et al.	5,949,977 A	9/1999	Hernandez
5,751,961 A	5/1998	Smyk	5,953,322 A	9/1999	Kimball
5,754,639 A	5/1998	Flockhart et al.	5,953,338 A	9/1999	Ma et al.
5,757,784 A	5/1998	Liebowitz et al.	5,953,392 A	9/1999	Rhie et al.
5,758,281 A	5/1998	Emery et al.	5,953,504 A	9/1999	Sokal et al.
5,761,290 A *	6/1998	Farris et al. 379/221.11	5,956,391 A	9/1999	Melen et al.
5,774,530 A	6/1998	Montgomery et al.	5,956,482 A	9/1999	Agraharam et al.
5,774,533 A	6/1998	Patel	5,956,697 A	9/1999	Usui
5,774,670 A	6/1998	Montulli	5,958,016 A	9/1999	Chang et al.
5,781,550 A	7/1998	Templin et al.	5,960,416 A	9/1999	Block
5,781,624 A	7/1998	Mitra et al.	5,966,427 A	10/1999	Shaffer et al.
5,784,617 A	7/1998	Greenstein et al.	5,966,434 A	10/1999	Schafer et al.
5,787,160 A	7/1998	Chaney et al.	5,970,065 A	10/1999	Miloslavsky
5,790,172 A	8/1998	Imanaka	5,978,773 A	11/1999	Hudetz et al.
5,790,806 A	8/1998	Koperda	5,978,840 A	11/1999	Nguyen et al.
5,793,672 A	8/1998	Wong et al.	5,987,108 A	11/1999	Jagadish et al.
5,796,790 A	8/1998	Brunner	5,995,503 A	11/1999	Crawley et al.
5,799,063 A	8/1998	Krane	5,995,606 A	11/1999	Civanlar et al.
5,799,317 A	8/1998	He et al.	5,999,524 A	12/1999	Corbalis et al.
5,802,146 A	9/1998	Dulman	6,011,975 A	1/2000	Emery et al.
5,802,502 A	9/1998	Gell et al.	6,014,379 A	1/2000	White et al.
5,802,510 A	9/1998	Jones	6,014,437 A	1/2000	Acker et al.
5,809,129 A	9/1998	Andersson et al.	6,014,687 A	1/2000	Watanabe et al.
5,812,656 A	9/1998	Garland et al.	6,016,307 A	1/2000	Kaplan et al.
5,812,776 A	9/1998	Gifford	6,018,567 A	1/2000	Dulman
5,812,795 A	9/1998	Horovitz et al.	6,021,126 A	2/2000	White et al.
5,812,865 A	9/1998	Theimer et al.	6,021,263 A	2/2000	Kujoory et al.
5,818,921 A	10/1998	Vander Meiden et al.	6,023,724 A	2/2000	Bhatia et al.

US 8,170,008 B2

Page 4

6,026,441 A	2/2000	Ronen	6,498,788 B1	12/2002	Emilsson et al.
6,029,203 A	2/2000	Bhatia et al.	6,519,254 B1	2/2003	Chuah et al.
6,035,028 A	3/2000	Ward et al.	6,529,516 B1	3/2003	Parzych
6,041,109 A	3/2000	Cardy et al.	6,556,659 B1	4/2003	Bowman-Amuah
6,041,117 A	3/2000	Androski et al.	6,584,093 B1	6/2003	Salama et al.
6,044,081 A	3/2000	Bell et al.	6,587,433 B1	7/2003	Borella et al.
6,047,051 A	4/2000	Ginzboorg et al.	6,594,254 B1	7/2003	Kelly
6,047,326 A	4/2000	Kilkki	6,600,733 B2	7/2003	Deng
6,049,545 A	4/2000	Stephenson et al.	6,636,485 B1	10/2003	Fijolek et al.
6,052,445 A	4/2000	Bashoura et al.	6,654,357 B1	11/2003	Wiedeman
6,052,450 A	4/2000	Allison et al.	6,671,285 B1	12/2003	Kirkby et al.
6,058,113 A	5/2000	Chang	6,678,264 B1	1/2004	Gibson
6,067,516 A	5/2000	Levay et al.	6,690,788 B1	2/2004	Bauer et al.
6,069,890 A	5/2000	White et al.	6,718,380 B1	4/2004	Mohaban et al.
6,073,160 A	6/2000	Grantham et al.	6,754,181 B1	6/2004	Elliott et al.
6,078,582 A	6/2000	Curry et al.	6,768,793 B1	7/2004	Sbisa et al.
6,084,956 A	7/2000	Turner et al.	6,775,519 B1	8/2004	Wiedeman et al.
6,088,358 A	7/2000	Tomita et al.	6,810,033 B2	10/2004	Derks
6,088,431 A	7/2000	LaDue	6,882,643 B1	4/2005	Mauger et al.
6,097,722 A	8/2000	Graham et al.	6,891,945 B2	5/2005	Porter et al.
6,097,804 A	8/2000	Gilbert et al.	6,983,040 B1	1/2006	Kalmanek, Jr. et al.
6,098,094 A	8/2000	Barnhouse et al.	7,013,001 B1	3/2006	Felger
6,101,182 A	8/2000	Sistanizadeh et al.	7,024,450 B1	4/2006	Deo et al.
6,104,704 A	8/2000	Buhler et al.	7,032,022 B1	4/2006	Shanungam et al.
6,108,314 A	8/2000	Jones et al.	7,061,923 B2	6/2006	Dugan et al.
6,108,704 A	8/2000	Hutton et al.	7,065,571 B2	6/2006	Schweitzer et al.
6,115,737 A	9/2000	Ely et al.	7,088,705 B2	8/2006	Curry et al.
6,125,113 A	9/2000	Farris et al.	7,170,887 B2	1/2007	Rosenberg
6,131,095 A	10/2000	Low et al.	7,185,073 B1	2/2007	Gai et al.
6,134,433 A	10/2000	Joong et al.	7,236,575 B2	6/2007	Kim et al.
6,137,777 A	10/2000	Vaid et al.	7,286,562 B1	10/2007	Vargo et al.
6,141,413 A	10/2000	Waldner et al.	2001/0039576 A1	11/2001	Kanada
6,141,686 A	10/2000	Jackowski et al.	2001/0040885 A1	11/2001	Jonas et al.
6,144,647 A	11/2000	Lopez-Torres	2002/0041590 A1	4/2002	Donovan
6,144,661 A	11/2000	Katsube et al.	2002/0064149 A1	5/2002	Elliott et al.
6,144,667 A	11/2000	Doshi et al.	2002/0083166 A1	6/2002	Dugan et al.
6,151,319 A	11/2000	Dommetry et al.	2002/0114324 A1	8/2002	Low et al.
6,154,777 A	11/2000	Ebrahim	2003/0039210 A1	2/2003	Jin et al.
6,157,648 A	12/2000	Voit et al.	2003/0193933 A1	10/2003	Jonas et al.
6,167,445 A	12/2000	Gai et al.	2003/0198335 A1	10/2003	Porter et al.
6,169,735 B1	1/2001	Allen, Jr. et al.	2004/0005046 A1	1/2004	Deo et al.
6,169,748 B1	1/2001	Barbas et al.	2004/0215787 A1	10/2004	Gibson et al.
6,175,618 B1	1/2001	Shah et al.	2005/0021713 A1	1/2005	Dugan et al.
6,181,690 B1	1/2001	Civanlar	2005/0191989 A1	9/2005	Plush et al.
6,181,703 B1	1/2001	Christie et al.	2006/0239254 A1	10/2006	Short et al.
6,185,198 B1	2/2001	LaDue	2007/0050492 A1	3/2007	Jorgensen
6,188,677 B1	2/2001	Oyama et al.	2007/0204036 A1	8/2007	Mohaban et al.
6,195,355 B1	2/2001	Demizu	2007/0263644 A1	11/2007	Christie et al.
6,198,738 B1	3/2001	Chang et al.			
6,205,148 B1	3/2001	Takahashi et al.			
6,212,162 B1	4/2001	Horlin			
6,233,604 B1	5/2001	Van Horne et al.			
6,240,091 B1	5/2001	Ginzboorg et al.			
6,243,373 B1	6/2001	Turock			
6,252,869 B1	6/2001	Silverman			
6,260,067 B1	7/2001	Barnhouse et al.			
6,278,693 B1	8/2001	Aldred et al.			
6,278,704 B1	8/2001	Creamer et al.			
6,282,281 B1	8/2001	Low			
6,298,120 B1	10/2001	Civanlar et al.			
6,298,383 B1	10/2001	Gutman et al.			
6,308,216 B1	10/2001	Goldschmidt et al.			
6,310,873 B1	10/2001	Rainis et al.			
6,324,264 B1	11/2001	Wiener et al.			
6,324,279 B1	11/2001	Kalmanek, Jr. et al.			
6,347,085 B2	2/2002	Kelly			
6,359,880 B1	3/2002	Curry et al.			
6,381,321 B1	4/2002	Brown et al.			
6,385,193 B1	5/2002	Civanlar et al.			
6,418,461 B1	7/2002	Barnhouse et al.			
6,418,467 B1	7/2002	Schweitzer et al.			
6,434,624 B1	8/2002	Gai et al.			
6,446,200 B1	9/2002	Ball et al.			
6,449,650 B1	9/2002	Westfall et al.			
6,456,617 B1	9/2002	Oda et al.			
6,463,470 B1	10/2002	Mohaban et al.			
6,466,570 B1	10/2002	Low et al.			
6,466,984 B1	10/2002	Naveh et al.			
6,470,081 B1	10/2002	Sbisa et al.			

FOREIGN PATENT DOCUMENTS

EP	0654930	5/1995
EP	0729281 A2	8/1996
EP	0740445	10/1996
EP	0767568	4/1997
EP	0781016	6/1997
EP	0802690	10/1997
EP	0812089	12/1997
EP	0823809	2/1998
EP	0729281 A3	1/2000
EP	1091528 A2	11/2001
EP	1091528 A3	11/2001
JP	09168051	6/1997
JP	09168063	6/1997
JP	09168064	6/1997
JP	09168065	6/1997
JP	09172459	6/1997
JP	09172462	6/1997
WO	9107839	5/1991
WO	9325035	12/1993
WO	9411813	5/1994
WO	9423523	10/1994
WO	9522221	8/1995
WO	9529564	11/1995
WO	9609714	3/1996
WO	9620553	4/1996
WO	9620448	7/1996
WO	9620553	7/1996
WO	9632800	10/1996
WO	9634341	10/1996

US 8,170,008 B2

Page 5

WO	9638018	11/1996
WO	9714238	4/1997
WO	9716007	5/1997
WO	9720424	6/1997
WO	9722211	6/1997
WO	9723078	6/1997
WO	9726749	7/1997
WO	9728628	8/1997
WO	9732427	9/1997
WO	9733412	9/1997
WO	9823080	5/1998
WO	9834391 A2	8/1998
WO	9834391 A3	8/1998

OTHER PUBLICATIONS

Mockapetris, "Domain Names—Concepts and Facilities", RFC 1034 (Nov. 1987).

Mockapetris, "Domain Names—Implementation and Specification", RFC 1035, 45 pages (Nov. 1987).

Malamud, et al., "Principles of Operation for the TPC.INT Subdomain: Remote Printing—Technical Procedures", RFC 1528, 10 pages (Oct. 1993).

Malamud, et al., "Principles of Operation for the TPC.INT Subdomain: Remote Printing—Administrative Policies", RFC 1529, 5 pages (Oct. 1993).

Malamud, et al., "Principles of Operation for the TPC.INT Subdomain: General Principles and Policy", RFC 1530, 6 pages (Oct. 1993).

Berners-Lee, "Universal Resource Identifiers in WWW", RFC 1630, 23 pages (Jun. 1994).

Manning, et al., "DNSNSAP Resource Records", RFC 1706, 9 pages (Oct. 1994).

"Web-on-Call Voice Browser Redefines Access to the Web", The Netphonic Press Release, (Mar. 1996).

Tao, et al., "Internet Access via Baseband and Broadband ISDN Gateways", IEEE International Conference on Computers and Communications, pp. 485-490 (Apr. 1994).

"Workstation Communications System", IBM Technical Disclosure Bulletin, vol. 37, No. 39 (Sep. 1994).

Bonetti, et al., "Distribution of RFC 1327 mapping rules via the Internet DNS: the INFNet distributed gateway system", Computer Networks and ISDN Systems, vol. 27, No. 3, pp. 461-469 (Dec. 1994).

Yang, et al., "The Design and Implementation of a Service Logic Execution Environment Platform", Globecom '93, vol. 3 (Nov. 1993).

Esaki, et al., "Abstraction and Control of Transport Network Resources for Intelligent Networks", Electronics & Communications in Japan: Part 1—Communications, vol. 76, No. 1 (Jan. 1993).

Chang, et al., "Rapid deployment of CPE-based telecommunications services", IEEE Globecom, vol. 2, pp. 876-880 (Nov. 1994).

Swale, et al., "Distributed Intelligence and data in public and private networks", BT Technology Journal, vol. 13, No. 2, pp. 95-104 (Apr. 1995).

Comer, et al., "Uniform Access to Internet Directory Services", Computer Communication Review, vol. 20, No. 4, pp. 50-59 (Sep. 1990).

Mach2 DNS (Domain Name Services) Module Au-A-11045/97 (Jun. 13, 1995).

Muffett, Proper Care and Feeding of Firewalls, Sun Microsystems, pp. 1-12 (Nov. 14, 1994).

Crocker, "RFC 822-Standard for the Format of ARPA Internet Text Messages", RFC 822, retrieved from <<http://www.faqs.org/rfcs/rfc822.html>>, pp. 1-42 (Aug. 13, 1982).

Girod, et al., "Requirements for URN Resolution Systems", retrieved from <<http://www.potaroo.net/ietf/ietf/draft-girod-urn-res-require/>>, pp. 1-31 (Jun. 13, 1996).

Hamilton, "Uniform Resource Identifiers & The Simple Discovery Protocol", Department of Computer Studies, Loughborough University of Technology, retrieved from <<http://martinh.net/uris/uris.html>>, 32 pages (Jun. 20, 1995).

Hardcastle-kille, "RFC 1279-X.500 and Domains", RFC 1279, retrieved from <<http://www.faqs.org/rfcs/rfc1279.html>>, pp. 1-12 (Nov. 1991).

Mealling, et al., "Uniform Resource Names, ISO OIDs and DNS", retrieved from <<http://www.gbiv.com/protocols/uri/drafts/draft-mealling-oid-dns-00.txt>>, pp. 1-14 (Nov. 22, 1995).

Mitra, et al., "Uniform Resource Names", retrieved from <<http://www.mitra.biz/uri/draft-ictrf-uri-resource-names-03.txt>>, pp. 1-9 (Nov. 1994).

Rose, et al., "RFC 1486—An Experiment in Remote Printing" RFC 1486, retrieved from <<http://www.faqs.org/rfcs/rfc1486.html>>, pp. 1-12 (Jul. 1993).

Shafer, et al., "URN Services", retrieved from <<http://staff.oclc.org/oclc/research/publications/shafer/urn/draft-shafer-uri-urn-resolution-00.html>>, pp. 1-10 (Jun. 1995).

Shafer, et al., "URN Services" retrieved from <<http://staff.oclc.org/oclc/research/publications/shafer/urn/draft-ietf-uri-urn-resolution-01.html>>, pp. 1-14 (Jul. 1995).

Sollins, et al., "RFC 1737—Functional Requirements for Uniform Resource Names", RFC 1737, retrieved from <<http://www.faqs.org/rfcs/rfc1737.html>>, pp. 1-6 (Dec. 1994).

The URN Implementors, "Uniform Resource Names", D-Lib Magazine, retrieved from <<http://www.dlib.org/dlib/February96/02arms.html>>, pp. 1-6 (Feb. 1996).

Everhart, et al., "RFC 1183—New DNS RR Definitions", retrieved from <<http://www.faqs.org/rfcs/rfc1183.html>>, pp. 1-9 (Oct. 1990).

Mockapetris, "RFC 1101—DNS Encoding of Network Names and Other Types", retrieved from <<http://www.faqs.org/rfcs/rfc1101.html>>, pp. 1-12 (Apr. 1989).

CT and the 'Net, "Audio and Video Over the Internet", Mar. 1996, 22 pages.

CT and the 'Net, "Supercharging the Web with Computer Telephony", Mar. 1996, 18 pages.

Ei-Gendy, H., et al., "Computer-Supported Routing for Intelligent Networks and Personalized Wired Communications", IEEE, Nov. 1994, pp. 1027-1033.

Grigoris, "Computer Telephony Over the Internet", CT and the Net, Mar. 1996, 6 pages.

McConnell, "How to Build an Internet PBX", <<http://www.phonezone.com/ip-phone.htm>>, Oct. 1996, pp. 1-9.

Sears "Innovations in Internet Telephony: The Internet as The Competitor to the POTS Network", Innovations in Internet Telephony: The Internet as the Successor to the POTS Network, Feb. 28, 1996, pp. 1-6.

Stylus Innovation, "Computer Telephony and the Internet", <<http://www.stylus.com/hvml.htm>>, Oct. 1996, 9 pages.

Welcome to the Phone Zone, Pacific Telephony Design, <<http://www.phonezone.com/index2.htm>>, obtained from Internet Oct. 28, 1996, pp. 1-9.

Yang, "INETPhone: Telephone Services and Servers on Internet". RFC 1789, Apr. 1995, Oct. 1996, pp. 1-6.

Ng, et al., "Distributed Architectures and Databases for Intelligent Personal Communication Networks", 1992 IEEE, ICWC '92, pp. 300-304, (Jun. 1992).

Malyan, et al., "Network Architecture and Signaling for Wireless Personal Communications", IEEE Journal on Selected Areas in Communications, vol. 11, No. 6, Aug. 1993, 12 pages.

Malyan, et al., "A Microcellular Interconnection Architecture for Personal Communications Networks", 1992 IEEE, pp. 502-505, (May 1992).

Cheriton, "Dissemination-Oriented Communication Systems: Final Report", Nov. 26, 1996, 18 pages.

Imilienski, et al., "Mobile Wireless Computing", Communications of the ACM, Oct. 1994, vol. 37, No. 10, pp. 18-28.

Bakre, et al., "M-RPC: A Remote Procedure Call Service for Mobile Clients" 1995 ACM 0-89791-814-S/95/10, pp. 97-110, (Nov. 1995).

O'Malley, et al., "A Dynamic Network Architecture", ACM Transactions on Computer Systems, vol. 10, No. 2, May 1992, pp. 110-143.

Cheshire, et al., "Internet Mobility 4x4", ACM 1996, <www.acm.org>, pp. 1-12, (Aug. 1996).

Chang et al., "Service Acquisition mechanism for the ClientService Model in Cygnus" IBM Canada Laboratory Technical Report TR 74.059, 1991, pp. 323-345, (May 1991).

Low, et al., "Webin—An Architecture for Fast Development of IN-Based Personal Services" Workshop Record Intelligent Network. Freedom and Flexibility: Realizing the Promise of Intelligent Network Services, IEEE, Apr. 21, 1996, 6 pages.

US 8,170,008 B2

Page 6

- Johnson, "Scalable Support for Transparent Mobile Host Internetworking", *Wireless Networks* 1 (1995), pp. 311-321, (Sep. 1995).
- Cobbold, et al., "Enhancements for Integrated Wireless Personal Communications Over Metropolitan Area Networks", *IEEE*, Apr. 1996, pp. 1370-1376.
- Jennings, "End-user Requirements for High-Integrity Directory", 1996, ICC '96, Conference Record, Converging Technologies for Tomorrow's Applications, 1996 IEEE International Conference on vol. 3, Jun. 23-27, 1996, pp. 1793-1796, vol. 3; Digital Object Identifier 10.1109/ICC.1996.535599.
- Anand, et al., "Accounting Architecture for Cellular Networks", 1996 *IEEE*, pp. 184-189, (Feb. 1996).
- Civanlar, et al., "FusionNet: Joining the Internet & Phone Networks for Multimedia Applications", 1996, pp. 431-432, (Nov. 1996).
- Karttunen, et al., "Cost Structure Analysis and Reference Model for SCALEABLE Network Services", *The Institution of Electrical Engineers*, 1996, pp. 3/1-3/9, (Nov. 1996).
- Chen, et al., "ATM and Satellite Distribution of Multimedia Educational Courseware", *IEEE*, Jun. 1996, pp. 1133-1137.
- Aidarous, et al., "The Role of the Element Management Layer in Network Management", *IEEE Network Operations and Management Symposium*, Feb. 14-17, 1994, pp. 59-70.
- Pontailier, "TMN and New Network Architectures", *IEEE Communications Magazine*, Apr. 1993, 6 pages.
- Li, et al., "Time-Driven Priority Flow Control for Real-Time Heterogeneous Internetworking", *Proceedings in Computer Communications, Fifteenth Annual Joint Conference of the IEEE Computer Communication Generation*, San Francisco, Mar. 24-28, 1996, vol. 1 Conf., Mar. 24, 1996 *IEEE*, pp. 189-197.
- Newton, "Telephony Messaging on the Internet", *Computer Telephony*, Jul. 1996, pp. 10-16.
- Gareiss, "Voice Over the Internet", *Data Communications*, vol. 25, No. 12, Sep. 1996, pp. 93, 94, 96, 98 and 100.
- Mills, "Phone Service Via the Internet May Slash Rates", *The Washington Post*, Aug. 11, 1996, p. A4.
- Venditto, "Internet Phones, The Future is Calling", *Internet World*, Jun. 1996, pp. 40-52.
- Carrel, "The TACACS+ Protocol, Version 1.75", *Network Working Group*, Lol Grant, Internet-Draft Systems, Oct. 1996, draft-grant-tacacs-00.txt, 64 pages.
- Low, "The Internet Telephony Red Herring", *IEEE* 1996, May 15, 1996, pp. 72-80.
- Chapman, et al., "Overall Concepts and Principles of TINA Version 1.0", *Telecommunications Information Networking Architecture Consortium*, TINA-C Deliverable. Date of Issue: Feb. 17, 1995, 65 pages.
- Parness, "IDT to Unveil Net2Phone At Comdex: Revolutionary System Uses Internet to Make Overseas PC-to-Telephone Calls for 10 cents a Minute", Nov. 6, 1995, <www.highbeam.com>, 2 pages.
- Fuente, et al., "Application of the TINA-C Management Architecture", 1995, 12 pages, (Apr. 1995).
- Finseth, "An Access Control Protocol, Sometimes Called TACACS", *University of Minnesota*, Jul. 1993 *Network Working Group*, Request for Comments: 1492, 22 pages.
- "VocalTec's Telephony Gateway—The Ultimate Internet Telephony Solution?", *Computer Telephony*, Sep. 1996, p. 30, 3 pages.
- Fang, "Building an Accounting Infrastructure for the Internet", *IEEE* Nov. 1996, pp. 105-109.
- Lipoff, "Operations Support System Framework for Data Over Cable Services", *Data Over Cable Technical Reports*, TR-DOCS-OS-SIW08-961016, Oct. 16, 1996, pp. 1-20.
- Edell, et al., abstract, "Billing Users and Pricing for TCP", *Journal on Selected Areas in Communications*, vol. 13, Issue 7, Sep. 1995, 2 pages.
- Duan et al., "Efficient Utilization of Multiple Channels Between Two Switches in ATM IN atm Networks", *IEEE Conference on Communications*, Jun. 1995, pp. 1906-1911.
- Ejiri, "For Whom the Advancing Service/Network Management", *IEEE Network Operations and Management Symposium*, Feb. 14-17, 1994, pp. 422-433.
- Elia, et al., "Skyplex: Distribution Up-Link for Digital Television via Satellite", *Apr. 1996*, pp. 305-313.
- Ely, "The Service Control Point as a Cross Network Integrator", *Apr. 1996*, pp. 1-8.
- Feinmann, "VIC Computer Telephony", *Computer Telephony*, Mar. 1996, 3 pages.
- Ganor, "Talk, Talk", *Tele.com*, Jun. 1996, pp. 68-72.
- Grami, et al., "The Role of Satellites in the Information Superhighway", *Jun. 1995*, pp. 1577-1581.
- Gupta, et al., "Technical Assessment of (T)INA-TMN-OSI Technology for Service Management Applications", *Feb. 1994*, pp. 877-887.
- Inamori, et al., "Applying TMN to a Distributed Communications Node System with Common Platform Software", *IEEE Feb. 1995*, pp. 83-87.
- Jain, "Evolving Existing Narrowband Networks Broadband Networks with IN Capabilities", *IEEE Apr. 1996*, 6 pages.
- Kim, "Talk is Cheap", *America's Network*, Jul. 15, 1996, pp. 34-39.
- Kishimoto, "Agent Communication System for Multimedia Communication Service", *IEEE Infocom*, Mar. 1996, *Proceedings vol. 1*, pp. 10-17.
- Fridisch, et al., "Terminals for Accessing the Internet—The Internet Telephone", *Alcatel Telecommunications Review—4th Quarter 1996*, pp. 304-309, (Dec. 1996).
- Diehl, "Data's New Voice", *BYTE*, Sep. 1996, pp. 129-135.
- Braun, et al., "A Framework for Flow-Based Accounting on the Internet", *Sep. 1993*, 5 pages.
- Margulies, "CT's Cyberdate with the 'Net'", *Aug. 1996*, *Computer Telephony Periscope*, pp. 28-29.
- Black, "OSI: A Model for Computer Communications Standards", *Prentice-Hall, Inc.*, 1991, pp. 157-201, (Oct. 1990).
- Arango, et al., *Guaranteed Internet Bandwidth*, *IEEE Nov. 1996*, pp. 862-866.
- Perret, et al., "MAP: Mobile Assistant Programming for Large Scale Communications Networks", *Apr. 1996*, pp. 1128-1132.
- Platt, "Why IsoEthernet Will Change the Voice and Video Worlds", *IEEE Communications Magazine*, Apr. 1996, pp. 55-59.
- Retkwa, "Telephone Politics", *Internet World*, Jun. 1996, pp. 54-60.
- Schreyer, et al., "Least Cost Call Routing", *IEEE Jun. 1996*, 7 pages.
- Serrano, "Evolution of Hybrid Fibre Coaxial Network for Multimedia Interactive Services", *British Telecommunications Engineering*, Oct. 1996, pp. 249-253.
- Sunaga, et al., "A Reliable Communication Switching Platform for Quick Provisioning", *IEEE Feb. 1995*, pp. 77-82.
- Telephony on the Internet, (Workshop Information) presented by International Quality & Productivity Center, IMTC and VOICE Technology & Services News, Sep. 1996, 6 pages.
- Tsuchida, et al., "Intelligent Dynamic Service Provisioning Architecture in the Multimedia Era", *Communications—Gateway to Globalization*, *IEEE Apr. 1996*, pp. 1117-1122.
- Estrin, et al., "Design Considerations for Usage Accounting and Feedback in Internetworks", 1990, 12 pages, (Oct. 1990).
- Lee, et al., "Integrated Packet Networks with Quality of Service Constraints", *Global Telecommunication Conference*, *IEEE Dec. 1991*, pp. 223-227.
- Eriksson, "MBONE: The Multicast Backbone", *Communications of the ACM*, vol. 37, No. 8, Aug. 1994, pp. 54-60, 13 pages.
- Gralla, "How the Internet Works", *Communication on the Internet*, Chapter 12, 1996, pp. 64-67, (Jun. 1996).
- Kahn, "Videoconferencing Debuts on the Internet", *Feb. 28, 1995*, 2 pages.
- Sullivan, "Videoconferencing Arrives on the Internet", *PC Week*, Aug. 1996, 1 page.
- Macedonia, et al., "Mbne Provides Audio and Video Across the Internet", *Apr. 1994*, 13 pages.
- Schulzrinne, "RFC 1890-RTP: Profile for Audio and Video Conferences with Minimal Control", *Jan. 1996*, 18 pages.
- Schulzrinne, "RFC 1889-RTP: A Transport Protocol for Real-Time Applications", *Jan. 1996*, 151 pages.
- Sclavos, et al., "Information model: From Abstraction to Application", *Feb. 1994*, 13 pages.
- Sriram, et al., "Voice Packerization and Compression in Broadband ATM Networks", *Apr. 1991*, *IEEE Journal on Selected Areas in Communications*, vol. 9, No. 3, pp. 294-304.
- SG15 Plenary, May 28, 1996, regarding H.323, 82 pages.
- RFC 1034, Nov. 1987, 41 pages.

US 8,170,008 B2

Page 7

- RFC 1035, Nov. 1987, 49 pages.
- SS7 Signaling Protocol, Message Transfer Part (MTP), Mar. 1996, 427 pages.
- SS7 Signaling Protocol, Signalling Connection Control Part (SCCP), Oct. 1992, 220 pages.
- Martersteck, et al., "The 5ESS Switching System: Introduction," AT&T Technical Journal, Jul. 1985, 10 pages.
- Kahane, "Introduction to VocalTec's CMA System, Intelligent Switching for a New Era of Telecom", Nov. 7, 1996, pp. 1-18.
- Kam, "Geek of the Week, Mobile IP Networking", Internet Talk Radio, Nov. 1993, 17 pages.
- Yamada, et al., "A Study on IN Basic Call State Model for Packet Switched Network", IEEE 1996, 18 pages, (Apr. 1996).
- ANSI, Signalling System No. 7 (SS7), General Information, ANSI, Jun. 1992, 78 pages.
- Cohen, Specifications for the Network Voice Protocol (NVP), Jan. 29, 1976, 33 pages.
- Everhart, et al., "New DNS RR Definitions", RFC 1183, Oct. 1990, 11 pages.
- Schulzrinne, "Simple Conference Invitation Protocol", Internet Draft, Feb. 22, 1996, 19 pages.
- Vocaltec's Telephony Gateway Software Captures 1996 Product of the Year Honors from Computer Telephony Magazine, Dec. 18, 1996, 2 pages.
- Verjinski, "PHASE, A Portable Host Access System Environment", IEEE 1989, 4 pages, (Oct. 1989).
- Schulzrinne, "Personal Mobility for Multimedia Services in the Internet", IDMS 1996, Mar. 1996, 18 pages.
- Garrahan, et al., "Intelligent Network Overview", IEEE Communications Magazine, Mar. 1993, 7 pages.
- Ash, et al., "Design & Optimization of Networks with Dynamic Routing", The Bell System Technical Journal, vol. 60, No. 8, Oct. 1981, pp. 1787-1820.
- Graff, et al., "The 5ESS Switching System: Physical Design/Hardware," AT&T Technical Journal, Jul. 1985, 46 pages.
- Bellcore Document GR-1428-CORE, CCS Network Interface Specification (CCSNIS), Supporting Toll-Free Service, A Module of CCSNIS, FR-905, May 1995, 93 pages.
- Email from Dr. Schulzrinne regarding publication of SCIP, Oct. 10, 1995, 1 page.
- GR-1280-CORE, Advanced Intelligent Network (AIN) Service Control Point (SCP) Generic Requirements, Telcordia Technologies Generic Requirements, Issue 1, Aug. 1993, 637 pages.
- GR-2863-CORE, CCS Network Interface Specification (CCSNIS) Supporting Advanced Intelligent Network (AIN), Bellcore, Generic Requirements, Issue 2, Dec. 1995, 148 pages.
- Office Action, U.S. Appl. No. 10/052,304 (U.S. Patent No. 7,206,304), mailed Jun. 19, 2006.
- Office Action, U.S. Appl. No. 11/066,880 (pending), mailed Apr. 14, 2009.
- Office Action, U.S. Appl. No. 11/066,880 (pending), mailed Nov. 23, 2009.
- Gibson, "Web Foes Belong in 'Animal House'," PC Week, Aug. 19, 1996, 1 page.
- VocalTech Homepage—2 pages, no earlier than Dec. 1996.
- Supplementary European Search Report for EP00976942 dated Jul. 7, 2003.
- International Search Report for PCT/US00/30448 dated Jan. 16, 2001, 3 pages.
- Haugk, et al., "The 5ESS Switching System: Maintenance Capabilities," AT&T Technical Journal, Jul. 1985, 32 pages.
- Basinger, et al., "The 5ESS Switching System: System Development Environment," AT&T Technical Journal, Jul. 1985, 18 pages.
- Bauer, et al., "The 5ESS Switching System: System Test, First-Office Application, and Early Field Experience," AT&T Technical Journal, Jul. 1985, 20 pages.
- Borum, et al., "The 5ESS Switching System: Hardware Design," AT&T Technical Journal, Jul. 1985, 21 pages.
- Byrne, et al., "The 5ESS Switching System: Applications Planning," AT&T Technical Journal, Jul. 1985, 23 pages.
- Carney, et al., "The 5ESS Switching System: Architectural Overview," AT&T Technical Journal, Jul. 1985, 18 pages.
- Delatore, et al., "The 5ESS Switching System: Factory System Testing," AT&T Technical Journal, Jul. 1985, 22 pages.
- Delatore, et al., "The 5ESS Switching System: Operational Software," AT&T Technical Journal, Jul. 1985, 28 pages.
- Fuhrer, et al., "The 5ESS Switching System: Operations, Administration, and Maintenance Capabilities," AT&T Technical Journal, Jul. 1985, 14 pages.
- Office Action, U.S. Appl. No. 11/590,489 (pending), mailed Mar. 18, 2009.
- Office Action, U.S. Appl. No. 11/590,489 (pending), mailed Dec. 1, 2009.
- Miller, "Managing the Internet", Troubleshooting TCP/IP, Chapter 7, 1992, pp. 365-375.
- Oppen, et al., "The Clearinghouse: A Decentralized Agent for Locating Named Objects in a Distributed Environment", ACM Transactions on Office Information Systems, vol. 1, No. 3, Jul. 1983, pp. 230-235.
- Sharp, et al., "Advanced Intelligent Networks—Now a Reality", Electronics & Communication Engineering Journal, Jun. 1994, pp. 153-162.
- Freiedes, et al., "ISDN opportunities for Large Business—800 Service Customers", IEEE International Conference on Communications, 1986, vol. 1, 7 pages, (Jun. 1986).
- Mark Handley, "Session Invitation Protocol," Feb. 22, 1996, 17 pages.
- Yeager, et al., "Web Server Technology, the Advanced Guide for World Wide Web Information Providers", Chapter 6: Searching for Information on the Web, 6.8.2. Using Database as Indexes to a Document Collection 1996, 3 pages, (Apr. 1996).
- Matta, et al., "Type-of-Service Routing in Dynamic Datagram Networks", 1994, Proceedings of the Conference on Computer Communications, IEEE, pp. 992-999, (Jun. 1994).
- Peeren, "The Intelligent Web", IEEE Apr. 1996, vol. 1, 13 pages.
- Clark, et al., "Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanism", 1992 ACM 0-89791-526-7/92/0008/0014, pp. 14-26, (Oct. 1992).
- Low, "The Internet Telephony Red Herring", Intelligent Networked Platforms Dept., Intelligent Networked Computing Laboratory, Hewlett-Packard Laboratories, Bristol, May 15, 1996, pp. 1-15.
- Waksberg, "Axe 10 and the Intelligent Network", Commutation and Transmission, No. 4, 1993, pp. 67-76, (Dec. 1993).
- Harmer, et al., "On the Move Project", Dec. 1996, pp. 1-44.
- Sisalem, et al., "The Network Video Terminal", IEEE 1996 Proceedings, pp. 3-10, (Aug. 1996).
- Kolarov, et al., "End-to-End Adaptive Rate Based Congestion Control Scheme for ABR Service in Wide Area ATM Networks", Feb. 1995, pp. 138-143.
- Niitsu, et al., "Computer-Aided Stepwise Service Creation Environment for Intelligent Network", Supercomm/International Conference on Communications, Jun. 14, 1992, pp. 454-458.
- Sairamesh, et al., "An Approach to Pricing, Optical Allocation and Quality of Service Provisioning in High Speed Packet Networks", Fourteenth Annual Joint conference of the IEEE Computer and Communication Societies, vol. 2, Apr. 1995, 21 pages.
- Sibal, et al., "Controlling Alternate Routing in General Mesh Packet Flow Networks", Association of Computing Machinery, 1994, pp. 168-179, (Aug. 1994).
- Morgan, et al., "Service Creation Technologies for the Intelligent Network", AT&T Technical Journal, 1991, pp. 58-71, (Jan. 1991).
- Fujioka, et al., "Universal Service Creation and Provision Environment for Intelligent Network", International Switching Symposium, May 1990, pp. 149-156.
- Moy, "RFC 1247—IETF Network Working Group: OSPF Version 2", Jul. 1991, <http://tools.ietf.org/html/rfc1247>, 183 pages.
- Ahmad, et al., "NBBS Traffic Management Overview, IBM Systems Journal, vol. 34, No. 4, 1995, pp. 640-648, (Dec. 1995).
- Braden, et al., "Integrated Services in the Internet Architecture: An Overview", RFC 1633, Jul. 1994, 28 pages.
- Bellamy, "Digital Telephony", John Wiley & Sons, Inc., 1982, pp. vii-ix, 9-14 and 238-240.
- Huitema, "Routing in the Internet", Prentice-Hall PTR, 1995, 168 pages, (Mar. 1995).
- Ziegler, "A Distributed Information System Study", IBM System Journal, vol. 18, No. 3, 1979, pp. 374-401, (Dec. 1979).

US 8,170,008 B2

Page 8

Elixman, et al., "OpenSwitching-Extending Control Architectures to Facilitate Applications", International Switching Symposium, vol. 2, Apr. 23, 1995, pp. 239-243.

Kabay, et al., "The Service Node—An Advanced IN Services Element", Technology Journal, vol. 13, No. 2, Apr. 1995, pp. 64-72.

Maruyama, "A Concurrent Object-Oriented Switching Program in Hill", IEEE Communications Magazine, vol. 29, No. 1, Jan. 1991, pp. 60-68.

Mayer, et al., "Service Net—2000: An Intelligent Network Evolution", AT&T Technical Journal, vol. 70, No. 3/4, 1991, pp. 99-110, (Jan. 1991).

Shabana, et al., "Intelligent Switch Architecture", Proceedings of the National Communications Forum, vol. 22, No. 2, Sep. 30, 1988, pp. 1312-1320.

Xiao, et al., abstract, "Sharing Model of Network Node Resources and Real-Time Scheduling", Journal: Mini-Micro Systems, vol. 16, No. 12, Dec. 1995, 1 page.

Gupta, et al., "Resource Sharing for Multi-Party Real-Time Communication", Conference Title: Proceedings IEEE INFOCOM '95, The Conference on Computer Communications, Fourteenth Annual Joint Conference of the IEEE Computer and Communications Societies, Bringing Information to people. Part vol. 3, 1995, pp. 1230-1237, (Apr. 1995).

Klein, et al., "Minimax Resource Allocation with Tree Structured Substitutable Resources", Journal: Operations Research, vol. 39, No. 2, Mar. 1991, pp. 285-295.

Weinrib, et al., abstract, "Decentralized Resource Allocation for Distributed Systems", Conference Title: IEEE INFOCOM '87, The Conference on Computer Communications. Proceedings. Sixth Annual Conference—Global Networks: Concept to Realization, Mar. 1987, 1 page.

* cited by examiner

U.S. Patent

May 1, 2012

Sheet 2 of 14

US 8,170,008 B2

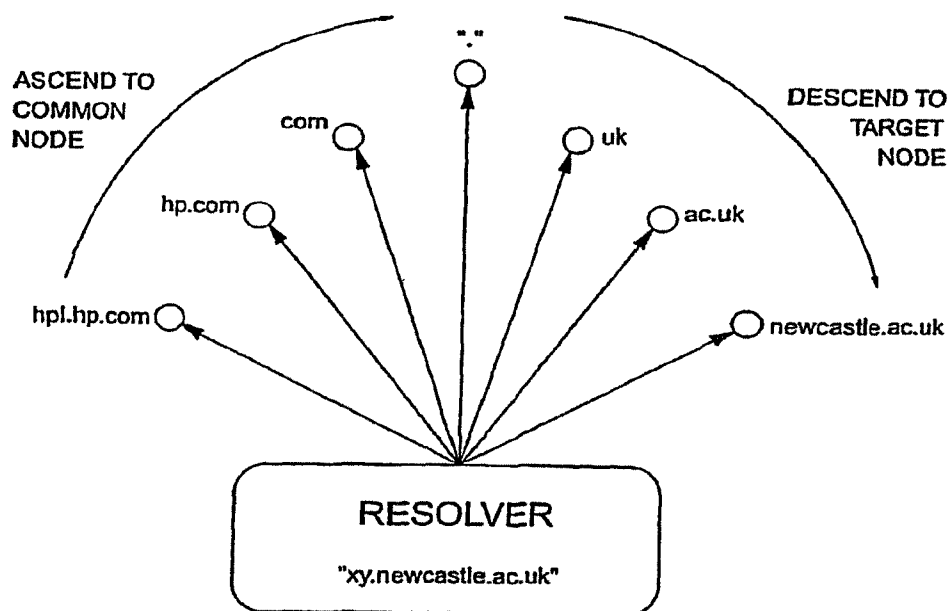



FIG. 3
(PRIOR ART)

U.S. Patent

May 1, 2012

Sheet 3 of 14

US 8,170,008 B2



http : //www.hp.com / Products.html

SCHEME HOST LOCATION ABSOLUTE PATH

FIG. 5
(PRIOR ART)




FIG. 6

U.S. Patent

May 1, 2012

Sheet 5 of 14

US 8,170,008 B2

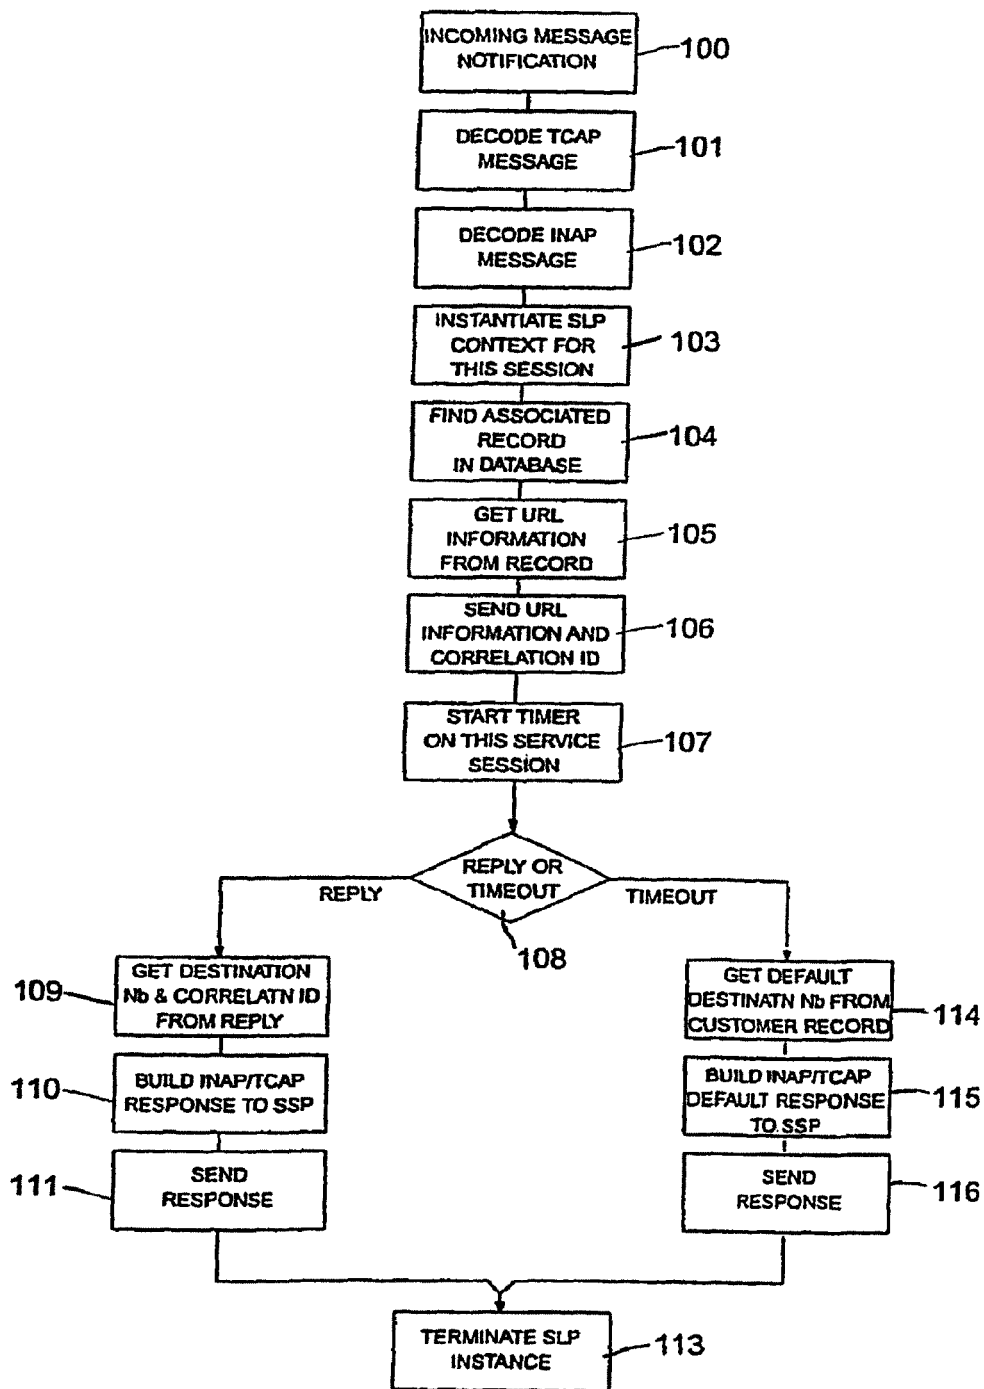


FIG. 7

Copy provided by USPTO from the PIRS Image Database on 03/06/2013

COM-DEL(0205)-00141737

A121

PTX002-0014

U.S. Patent

May 1, 2012

Sheet 6 of 14

US 8,170,008 B2




FIG. 8

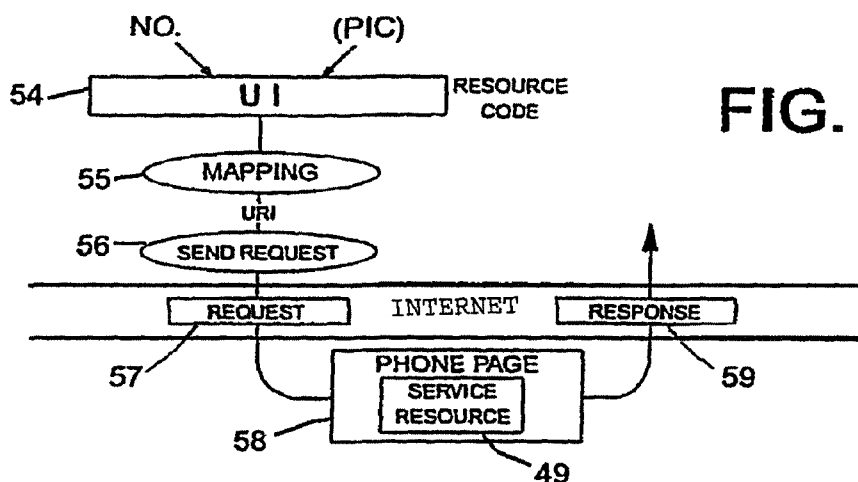


FIG. 9




FIG. 10

U.S. Patent

May 1, 2012

Sheet 7 of 14

US 8,170,008 B2




FIG. 11

U.S. Patent

May 1, 2012

Sheet 8 of 14

US 8,170,008 B2




FIG. 12A




FIG. 12B




FIG. 12C

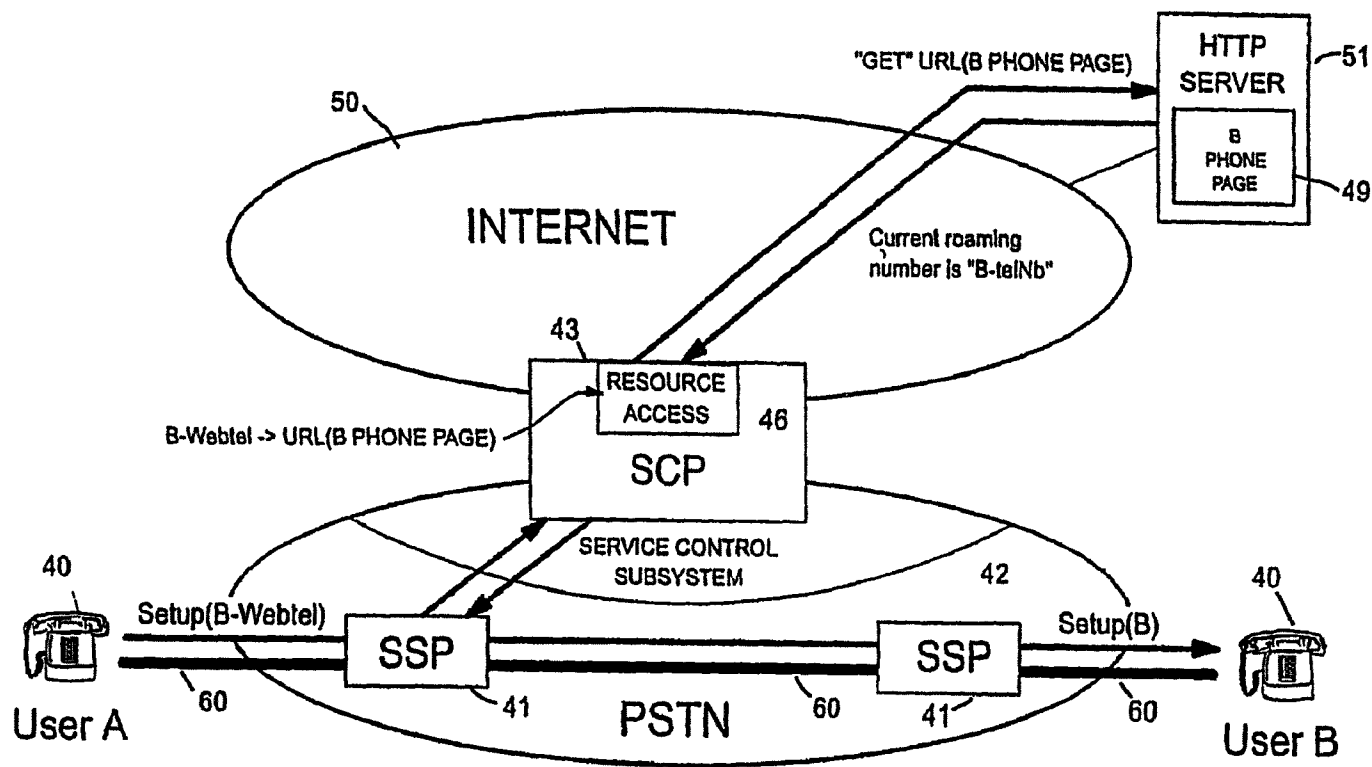


FIG. 13

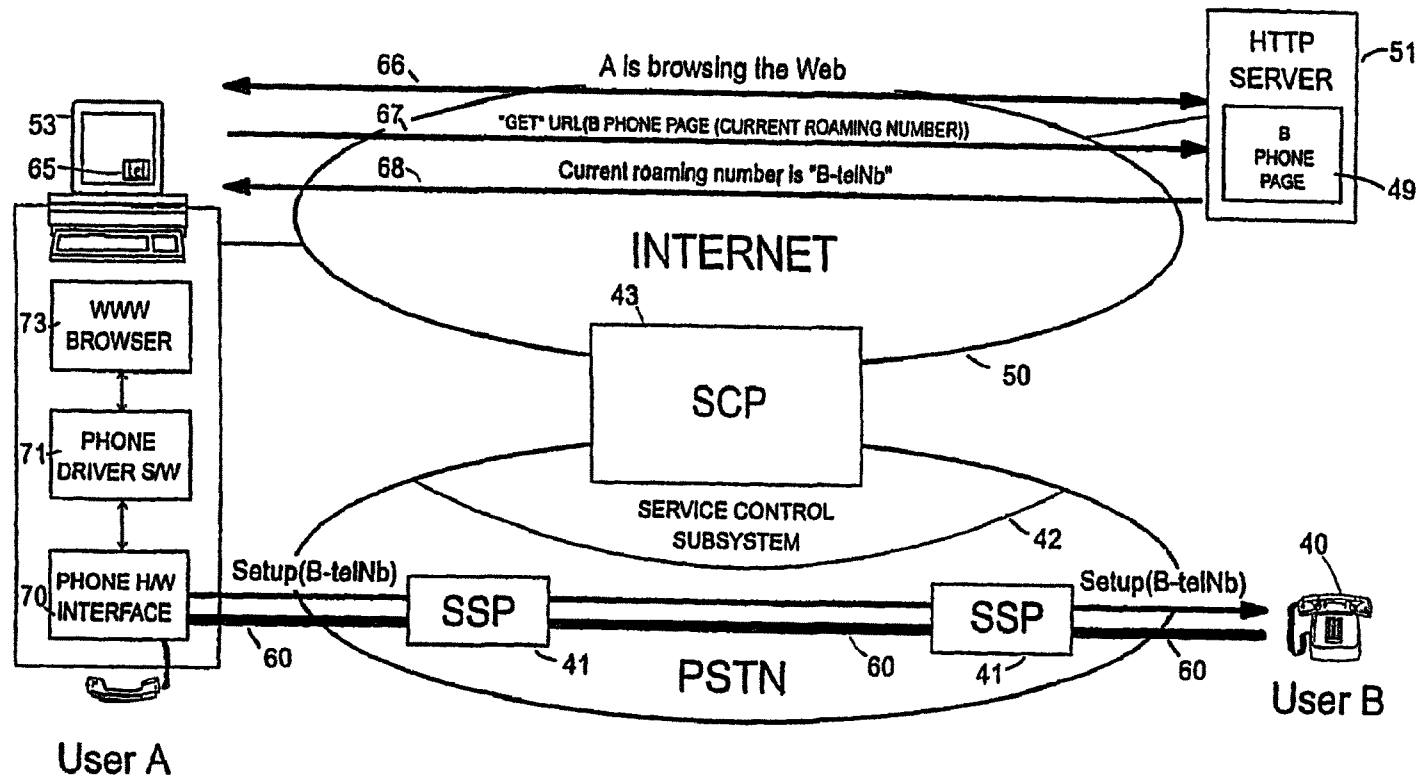


FIG. 14

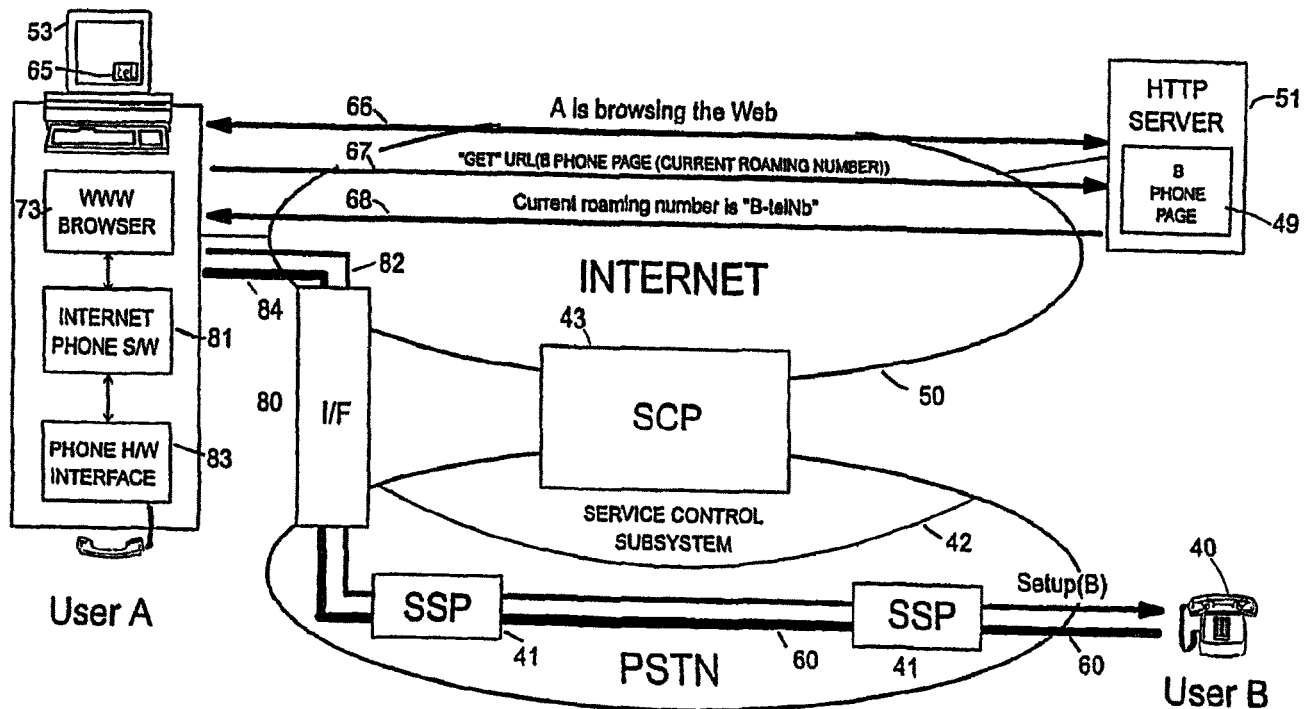


FIG. 15

U.S. Patent

May 1, 2012

Sheet 12 of 14

US 8,170,008 B2

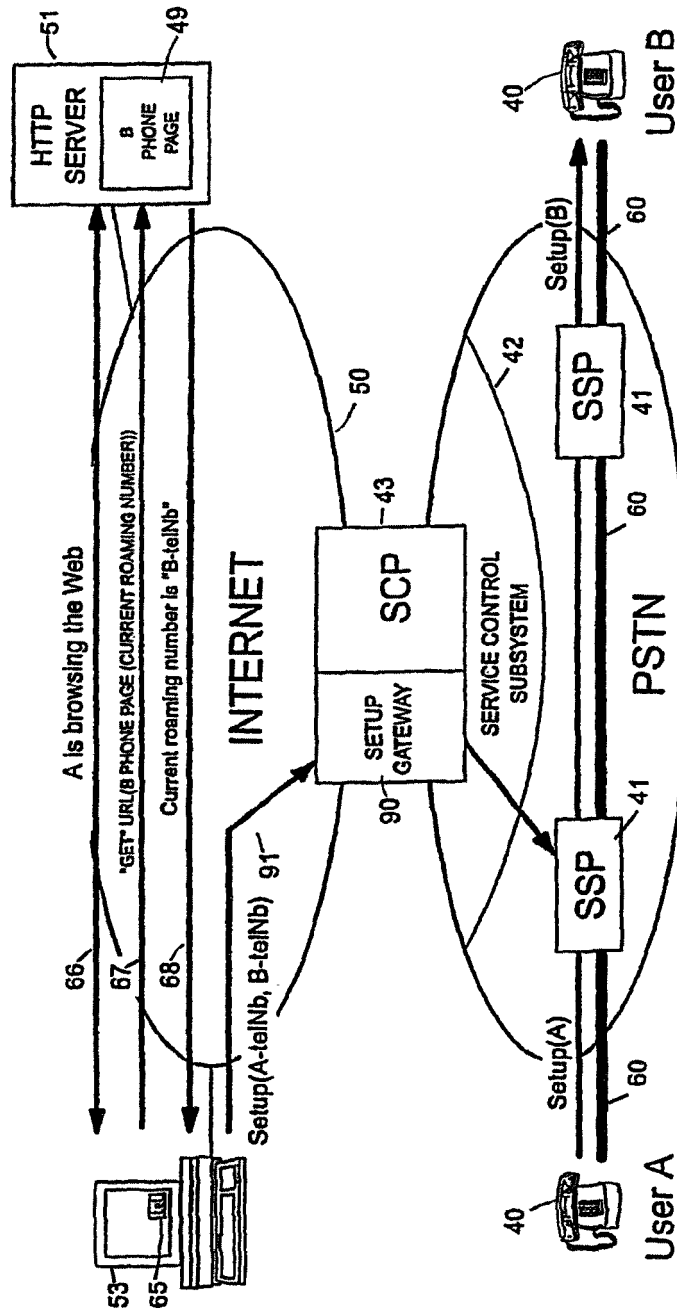


FIG. 16

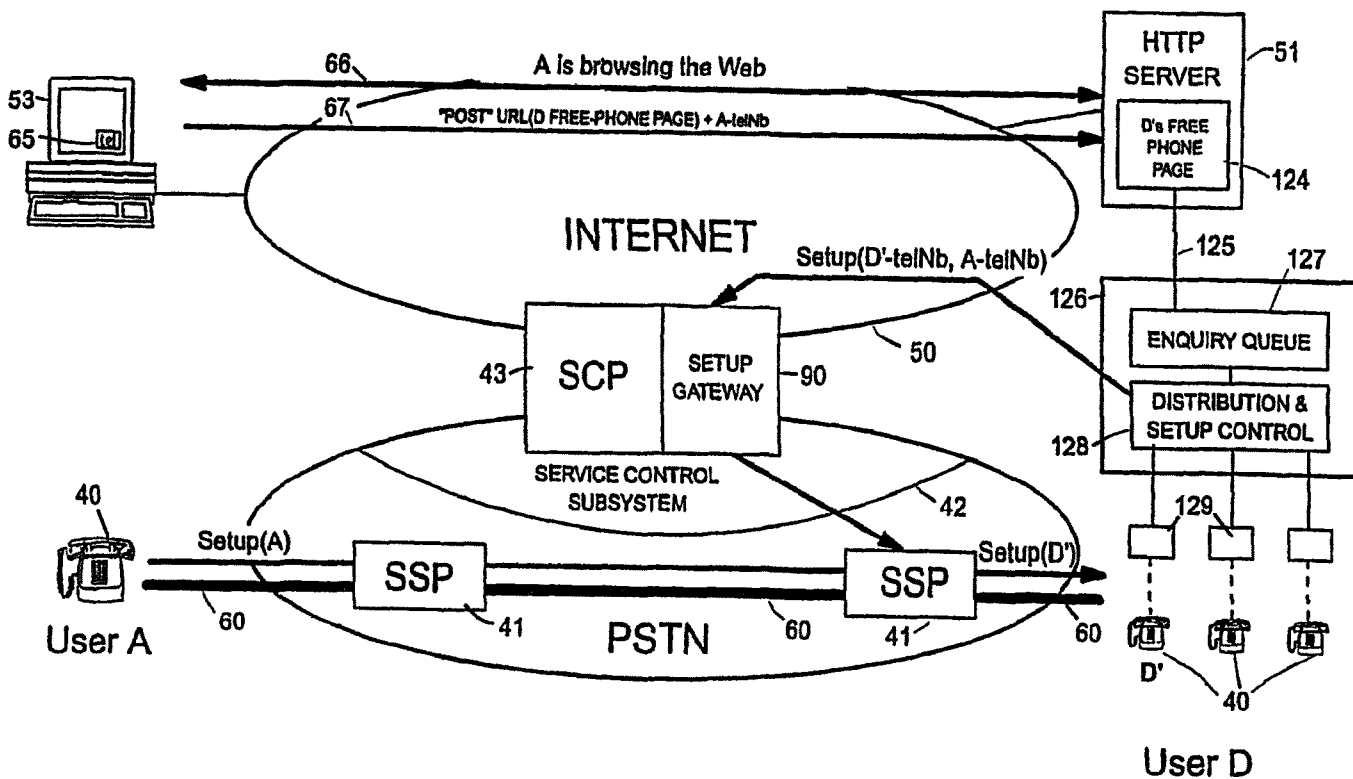


FIG. 17

U.S. Patent

May 1, 2012

Sheet 14 of 14

US 8,170,008 B2




FIG. 18

US 8,170,008 B2

1

METHOD AND APPARATUS FOR ACCESSING COMMUNICATION DATA RELEVANT TO A TARGET ENTITY IDENTIFIED BY A NUMBER STRING

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application is a continuation of U.S. patent application Ser. No. 11/066,880, filed Feb. 25, 2005, which is a divisional of U.S. patent application Ser. No. 10/052,285, filed Jan. 18, 2002 (now U.S. Pat. No. 7,012,916), which is a divisional of U.S. patent application Ser. No. 09/077,795, filed Jun. 5, 1998 (now U.S. Pat. No. 6,466,570), which claims priority under 35 USC 371 to PCT application No. PCT/GB96/03055, filed Dec. 11, 1996, which further claims priority to Great Britain Patent Application No. GB 9603582.9, filed Feb. 20, 1996, to European Patent Application No. EP 95410148.1, filed Dec. 22, 1995, and to Great Britain Patent Application No. GB 9525190.6, filed Dec. 11, 1995. All of the above-listed patent applications, publications thereof, and patents, are hereby incorporated by reference herein as to their entireties.

FIELD OF THE INVENTION

The present invention relates to a method of accessing service resource items that are intended to be used in setting up bearer channels through a switched telecommunications system.

As used herein, the term "switched telecommunication system" means a system comprising a bearer network with switches for setting up a bearer channel through the network. The term "switched telecommunication system" is to be taken to include not only the existing public and private telephone systems (whether using analogue phones or ISDN-based), but also broadband (ATM) and other switch-based bearer networks that are currently being implemented or may emerge in the future. For convenience, the term "switched telecommunication system" is sometimes shortened herein to telecommunication system.

Reference to a "call" in the context of a switched telecommunication system is to be understood as meaning a communication through a bearer channel set up across the bearer network, whilst references to call setup, maintenance and takedown are to be taken to mean the processes of setting up, maintaining and taking down a bearer channel through the bearer network. Terms such as "call processing" and "call handling" are to be similarly interpreted.

The term "communication system" when used herein should be understood as having a broader meaning than switched telecommunication system, and is intended to include datagram-based communication systems where each data packet is independently routed through a bearer network without following a predetermined bearer channel.

BACKGROUND OF THE INVENTION

Telecommunication companies running PSTNs (Public Switched Telephone Networks) and PLMNs (Public Land Mobile Networks) are in the business of providing communication services and in doing so are providing increasing built-in intelligence in the form of "IN services" such as 800 number services and call forwarding. In contrast, the World Wide Web (WWW), which has seen explosive growth in recent times, is an example of an Internet-based global network providing complex information services. These two

2

worlds, that of the large communications utilities and that of the highly dynamic, pioneer-spirit WWW information culture, are uneasy companions and each plans to encroach on the domain previously occupied by the others; thus telephony services will be offered over the WWW and information services over the public communication infrastructure.

The present invention proposes technologies for a more synergetic relationship between these two worlds than is currently envisaged and in order to place the present invention in context, a review will first be given of each of these two worlds.

Telephone Networks with IN Services

The Basic PSTN. The basic service provided by a PSTN (Public Switched Telephone Network) is the interconnection of two telephones (that is, setting up a bearer channel between the telephones) according to a called-party telephone number input at the calling-party telephone. FIG. 1 is a simplified representation of a PSTN providing such a service. In particular, customer premises equipment, CPE, 10 (such as standard analogue telephones, but also more recently ISDN terminals) are connected through an access network 11 to switching points, SPs 12. The SPs 12 form nodes in an inter-exchange network 13 made up of interconnecting trunks 14 and SPs that are controlled by control entities 15 in the SPs.

The control effected by the control entities 15 is determined by signalling inputs received from the CPEs and other SPs, and involves call setup, maintenance and clearance to provide the desired bearer channel between calling CPE and called CPE. Conceptually, the PSTN may be thought of as a bearer network and a control (signalling) network, the function of the latter being to effect call control through the bearer network, namely the control of setup, maintenance and take down of bearer channels through the bearer network; in practice, the bearer and signalling networks may use the same physical circuits and even the same logical channels.

Thus, where the CPE is a traditional dumb telephone, control signalling between the CPE and its local SP is in-band signalling, that is, the signalling is carried on the same channel as used for voice; this signalling is interpreted and converted at the SPs 12 into signalling between SPs that uses a dedicated common-channel signalling network 16 (implemented nowadays using the SS7 protocol suite). Where the CPE is an ISDN terminal, signalling is carried in a separate channel directly from the CPE on an end-to-end. Modern SPs use the ISUP (ISDN User Part) SS7 protocol for inter-exchange call control signalling whether the CPE is a standard telephone or an ISDN terminal.

Telephone Numbering Plans—As certain aspects of the present invention are influenced by the structuring of telephone numbers, a brief description will now be given of the structuring of such numbers. Telephone numbers form an international, hierarchical addressing scheme based on groups of decimal digits. The top level of the hierarchy is administered by the ITU-T, which has allocated single-digit numeric codes to the major geographic zones (for example "1" for North America, "2" for Africa, "3" for Europe, "4" for Europe, "5" for South America and Cuba, etc.). Within each zone countries are assigned 2 or 3 digit codes, so that within zone 3 France is "33", and within zone 4 the UK is "44". Administration of the numbering plan within a country is delegated to a national body, such as the Office of Telecommunications ("OfTel") in the UK. The following further description is based on the UK numbering plan, but the scheme described will be recognised as having widespread applicability.

In the UK all national numbers are prefixed by a code from 01 to 09 (the '0' prefix is dropped in international dialing).

US 8,170,008 B2

3

The currently assigned codes are "01" for Geographic Area Codes, "02" for Additional Geographic Area Codes, "04" for Mobile Services, "07" for Personal Numbers, and "08" for Special Service (freephone, information). Normal wireline PSTN subscriber telephone numbers are allocated from the Geographic Area Code codes, and currently only codes prefixed by 01 are allocated. Geographic area codes are presently 3 or 4 digits (excluding the leading '0') and there are currently 638 geographic areas each with its own code. A full national UK dialed number takes two forms:

0	171	634 8700
	area code	local number (7 digit)
0	1447	456 987
	area code	local number (6 digit)

The first case has the '0' prefix, a 3 digit area code and a 7 digit local number, and the second case has the '0' prefix, a 4 digit area code, and a 6 digit local number. Further interpretation of the local number will take place within the area exchange, as even a 6 digit address space is too large for a single switch, and for a typical local area several switches may be needed to host the required number of subscriber lines. This interpretation is opaque and is a matter for the area service provider.

In the current PSTN the inherently hierarchical and geographic interpretation of telephone numbers is mirrored by the physical architecture of the network. A telephone number is structured in a way that makes it easy to route a call through the network. At each step, the prefix of the number provides information about the current routing step, and the suffix (perhaps opaquely) provides information about subsequent routing steps; as long as a switch knows how to parse a prefix and carry out a routing step, it does not need to understand the content of the suffix, which is left for subsequent routing steps. For this reason the international and national switching fabric is also organised hierarchically.

Intelligent Networks. Returning now to a consideration of the current telephone network infrastructure, in addition to basic call handling, an SP may also serve to provide what are called IN (Intelligent Network) services; in this case the SP is termed a service switching point, SSP. An SSP 25 is arranged to suspend call processing at defined points-in-call upon particular criteria being met, and to delegate the continuation of call processing to a service control subsystem providing a service control function (SCF) either in the form of a service control point, SCP 17 (see FIG. 2) or an Adjunct 18. The Adjunct 18 is directly associated with an SSP 25 whilst the SCP 17 and SSP 25 communicate with each other via an extended common channel signalling (CCS) network 16 that may include signal transfer points (STP) 19. The SCP 17 may be associated with more than one SSP 25. Both the SCP 17 and Adjunct 18 provide a service logic execution environment (SLEE) 20 in which instances of one or more service logic programs (SLP) 21 can execute. The SLEE 20 and SLP 21 together provide service control functionality for providing services to the SSP 25.

Service logic running in an SCP or Adjunct will generally make use of subscriber information stored in a service data function (SDF) 22 that may be integral with the SCP/Adjunct or partially or wholly separate therefrom. The service data function (SDF), like the service control function (SCF) forms part of the service control subsystem of the PSTN. It may be noted that some or all of the service control function may be built into the PSTN switches themselves.

4

In addition to the SCP 17 and Adjunct 18, the FIG. 2 network includes an intelligent peripheral (IP) 23. The IP 23 provides resources to the SSP 25 such as voice announcements and DTMF digit collection capabilities. The network will also include an operation system (not shown) that has a general view of the network and its services and performs functions such as network monitoring and control.

In operation, when the SSP 25 receives a call, it examines internal trigger conditions and, possibly, user information (eg dialed digits) to ascertain if the call requires a service to be provided by the service control subsystem 17, 18; the checking of trigger conditions may be carried out at several different points in call processing. Where the SSP 25 determines that a service is required it messages the service control subsystem (either SCP 17 or Adjunct 18) requesting the desired service and sending it a logic representation of the call in terms of its connectivity and call processing status. The service control subsystem then provides the requested service and this may involve either a single interaction between the SSP and service control subsystem or a session of interactions. A typical service is call forwarding which is a called-party service giving expression to an end-user requirement as simple as "if you call me on number X and it rings ten times, try calling number Y". In this case, it is the SSP local to the called end-user that triggers its associated SCP (or Adjunct) to provide this service; it will, of course, be appreciated that the SSP must be primed to know that the service is to be provided for a called number X.

The above-described model for the provision of IN services in a PSTN can also be mapped onto PLMNs (Public Land Mobile Networks) such as GSM and other mobile networks. Control signalling in the case of a mobile subscriber is more complex because in addition to all the usual signalling requirements, there is also a need to establish where a call to a mobile subscriber should be routed; however, this is not a very different problem from a number of called-party IN services in the PSTN. Thus in GSM, the service-data function (SDF) is largely located in a system named a Home Location Register (HLR) and the service control function in a system named a Visitor Location Register (VLR) that is generally associated on a one-to-one basis with each SSP (which in GSM terminology is called a Mobile Switching Centre, MSC).

Because subscribers are mobile, the subscriber profile is transported from the HLR to whichever VLR happens to be functionally closest to be mobile subscriber, and from there the VLR operates the (fixed) service using the subscriber profile and interacts with the SSP. The HLR and VLR thus constitute a service control subsystem similar to an SCP or Adjunct with their associated databases.

It is, of course, also possible to provide IN services in private telephone systems and, in this case, the service control function and service data function are generally either integrated into a PABX (Private Automatic Branch Exchange) or provided by a local computer. The service control subsystem, whilst present, may thus not be a physically distinct from the PABX.

The above-described general architectural framework for providing IN services has both strengths and flaws. Its main strength is that it works and many services have been successfully deployed, such as 800 number services, credit card calling, voicemail, and various call waiting and redirection services. However, despite years of standardisation, services are still implemented one-at-a-time on proprietary platforms and do not scale well. The approach has been based on large, fault-tolerant systems which provide services for hundreds of thousands or even millions of subscribers and take years to

US 8,170,008 B2

5

deploy. Furthermore, since the networks used to support these services also constitute the basic telephone infrastructure, anything attached to these networks must be rigorously vetted. Additionally, each country and operator tends to have local variations of the so-called standards making it difficult to supply standard products and thereby braking the dynamics of competition.

The World Wide Web

In contrast to the slow deliberate progress of the telephone infrastructure, the WWW has grown explosively from its inception in 1989 to become the primary electronic information distribution service in terms of spread, availability and richness of information content. Anyone can, for a modest outlay, become an information provider with a world-wide audience in a highly interconnected information architecture.

The WWW is a client-server application running over the Internet and using a client-server protocol which mandates only the simplest of exchanges between client and server. This protocol is HTTP (Hyper Text Transfer Protocol) which is optimised for use over TCP/IP networks such as the Internet; the HTTP protocol can, however, also be used over networks using different communication protocol stacks.

Since the availability of literature concerning the WWW has seen the same sort of growth as the WWW itself, a detailed description of the WWW, HTTP and the Internet will not be given herein. An outline description will, however, be given with attention being paid to certain features of relevance to the present invention.

The WWW uses the Internet for interconnectivity. Internet is a system that connects together networks on a worldwide basis. Internet is based on the TCP/IP protocol suite and provides connectivity to networks that also use TCP/IP. For an entity to have a presence on the Internet, it needs both access to a network connected to the Internet and an IP address. IP addresses are hierarchically structured. Generally an entity will be identified at the user level by a name that can be resolved into the corresponding IP address by the Domain Name System (DNS) of the Internet. Because the DNS or adaptations of it are fundamental to at least certain embodiments of the invention described hereinafter, a description will next be given of the general form and operation of the DNS.

The Domain Name System—The DNS is a global, distributed, database, and without its performance, resilience and scalability much of the Internet would not exist in its current form. The DNS, in response to a client request, serves to associate an Internet host domain name with one or more Registration Records (RR) of differing types, the most common being an address (A) record (such as 15.144.8.69) and mail exchanger (MX) records (used to identify a domain host configured to accept electronic mail for a domain). The RRs are distributed across DNS name servers world-wide, these servers cooperating to provide the domain name translation service; no single DNS server contains more than a small part of the global database, but each server knows how to locate DNS servers which are “closer” to the data than it is. For present purposes, the main characteristics of the DNS of interest are:

The host name space is organised as a tree-structured hierarchy of nodes with each host having a corresponding leaf node; each node has a label (except the root node) and each label begins with an alphabetic character and is followed by a sequence of alphabetic characters or digits. The full, or “fully qualified” name of a host is the string of node labels, each separated by a “.”, from the corresponding leaf node to the root node of the hierarchy, this latter being represented by a terminating “.” in the response. Thus a host machine “fred” of

6

Hewlett-Packard Laboratories in Bristol, England will have a fully qualified domain name of “fred.hpl.hp.com.” (note that if a host name does not have a terminal “.” it is interpreted relative to the current node of the naming hierarchy).

Each host has one or more associated Registration Records (RRs).

There are a plurality of DNS servers each with responsibility for a subtree of the name space. A DNS server will hold RRs for all or part of its subtree—in the latter case it delegates responsibility for the remainder of the subtree to one or more further DNS servers. A DNS server knows the address of any server to which it has delegated responsibility and also the address of the server which has given it the responsibility for the subtree it manages. The DNS servers thus point to each other in a structuring reflecting that of the naming hierarchy.

An application wishing to make use of the DNS does so through an associated “resolver” that knows the address of at least one DNS server. When a DNS server is asked by this resolver for an RR of a specified host, it will return either the requested RR or the address of a DNS server closer to the server holding the RR in terms of traversal of the naming hierarchy. In effect, the hierarchy of the servers is ascended until a server is reached that also has responsibility for the domain name to be resolved; thereafter, the DNS server hierarchy is descended down to the server holding the RR for the domain name to be resolved.

The DNS uses a predetermined message format (in fact, it is the same for query and response) and uses the IP protocols.

These characteristics of the DNS may be considered as defining a “DNS-type” system always allowing for minor variations such as in label syntax, how the labels are combined (ordering, separators), the message format details, evolutions of the IP protocols etc.

Due to the hierarchical naming structure, it is possible to delegate responsibility for administering domains (subtrees) of the name space recursively. Thus, the top-level domains are administered by InterNic (these top-level domains include the familiar ‘com’, ‘edu’, ‘org’, ‘int’, ‘net’, ‘mil’ domains as well as top-level country domains specified by standard two-letter codes such as ‘us’, ‘uk’, ‘fr’ etc.). At the next level, by way of example Hewlett-Packard Company is responsible for all names ending in ‘hp.com’ and British Universities are collectively responsible for all names ending in ‘ac.uk’. Descending further, and again by way of example, administration of the domain ‘hpl.hp.com’ is the responsibility of Hewlett-Packard Laboratories and administration of the subtree (domain) ‘newcastle.ac.uk’ is the responsibility of the University of Newcastle-upon-Tyne.

FIG. 3 illustrates the progress of an example query made from within Hewlett-Packard Laboratories. The host domain name to be resolved is ‘xy.newcastle.ac.uk’, a hypothetical machine at the University of Newcastle, United Kingdom. The query is presented to the DNS server responsible for the ‘hpl.hp.com’ subtree. This server does not hold the requested RR and so responds with the address of the ‘hp.com’ DNS server; this server is then queried and responds with the address of the ‘com’ DNS server which in turn responds with the address of the ‘.’ (root) DNS server. The query then proceeds iteratively down the ‘uk’ branch until the ‘newcastle.ac.uk’ server responds with the RR record for the name ‘xy’ in its subtree.

This looks extremely inefficient, but DNS servers are designed to build a dynamic cache, and are initialised with the addresses of several root servers, so in practice most of the iterative queries never take place. In this case the ‘hpl.hp.com’ DNS server will know the addresses of several root servers, and will likely have the addresses of ‘uk’ and ‘ac.uk’

US 8,170,008 B2

7

servers in its cache. The first query to the 'hpl.hp.com' server will return the address of the 'ac.uk' server. The second query to the 'ac.uk' server will return the address of the 'newcastle.ac.uk' server, and the third query will return the RR in question. Any future queries with a 'newcastle.ac.uk' prefix will go direct to the newcastle DNS server as that address will be retained in the 'hpl.hp.com' DNS server cache. In practice names within a local subtree are resolved in a single query, and names outside the local subtree are resolved in two or three queries.

Rather than a resolver being responsible for carrying out the series of query iterations required to resolve a domain name, the resolver may specify its first query to be recursive in which case the receiving DNS server is responsible for resolving the query (if it cannot directly return the requested RR, it will itself issue a recursive query to a 'closer' DNS server, and so on).

It should also be noted that in practice each DNS server will be replicated, that is, organised as a primary and one or more secondaries. A primary DNS nameserver initialises itself from a database maintained on a local file system, while a secondary initialises itself by transferring information from a primary. A subtree will normally have one primary nameserver and anything up to ten secondaries—the limitation tends to be the time required by the secondaries to update their databases from the primary. The primary database is the master source of subtree information and is maintained by the domain DNS administrator. The secondaries are not simply standby secondaries but each actively participates in the DNS with dependent servers that point to it rather than to the corresponding primary.

DNS implementations, such as BIND, are widely available as a standard part of most UNIX systems, and can claim to be among the most robust and widely used distributed applications in existence.

Operation of the WWW Referring now to FIG. 4 of the accompanying drawings, access to the Internet 30 may be by direct connection to a network that is itself directly or indirectly connected to the Internet; such an arrangement is represented by terminal 31 in FIG. 4 (this terminal may, for example, be a Unix workstation or a PC). Having a connection to the Internet of this form is known as having 'network access'. Any entity that has network access to the Internet may act as a server on the Internet provided it has sufficient associated functionality; in FIG. 4, entity 32 with file store 37 acts as a server.

Many users of the WWW do not have network access to the Internet but instead access the Internet via an Internet service provider, ISP, 33 that does have network access. In this case, the user terminal 34 will generally communicate with the ISP 33 over the public telephone system using a modem and employing either SLIP (Serial Line Interface Protocol) or PPP (Point-to-Point Protocol). These protocols allow Internet packets to traverse ordinary telephone lines. Access to the Internet of this form is known as "dialup IP" access. With this access method, the user terminal 34 is temporarily allocated an IP address during each user session; however, since this IP address may differ between sessions, it is not practical for the entity 34 to act as a server.

A cornerstone of the WWW is its ability to address particular information resources by means of an Uniform Resource Identifier (URI) that will generally be either a Uniform Resource Locator (URL) that identifies a resource by location, or a Uniform Resource Name (URN) that can be resolved into an URL. By way of example, a full or "absolute" URL will comprise the following elements:

8

scheme	this is the access scheme to be used to access the resource of interest;
host	the Internet host domain name or IP address;
port	the host port for the (TCP) connection;
abs-path	the absolute path of the resource on the host.

In fact, the 'port' may be omitted in which case port 80 is assumed.

FIG. 5 of the accompanying drawings shows an example URL for the Hewlett-Packard products welcome page. In this case, the elements are:

scheme	http
host	www.hp.com
port	omitted (port 80 assumed)
abs-path	Products.html

The HTTP protocol is based on a request/response paradigm. Referring again to FIG. 4 of the drawings, given a particular URI identifying a resource 30 to be accessed, a client establishes a connection with the server 31 corresponding to the "host" element of the URI and sends a request to the server. This request includes a request method, and the "Request-URI" (which is generally just the absolute path of the resource on the server as identified by the "abs-path" element of the URI); the request may include additional data elements. The server 31 then accesses the resource 36 (here held on storage 37) and responds and this response may include an entity of a type identified by a MIME (Multipurpose Internet Mail Extensions) type also included in the response.

The two main request methods are:

GET—This method results in the retrieval of whatever information (in the form of an entity) is identified by the Request-URI. It is important to note that if the Request-URI refers to a data-producing process, it is the produced data which is returned as the entity in the response and not the source text of the process.

POST—This method is used to request that the destination server accept the entity enclosed in the request as a new subordinate of the resource identified by the Request-URI. The POST method can be used for annotation of existing resources, providing a message to a bulletin board, providing data to a data-handling process (for example, data produced as the result of submitting a form), and extending a database through an append operation.

In summary, the GET method can be used to directly retrieve data, or to trigger any process that will return an entity (which may either be data or a simply an indication of the result of running the process). The POST method is used for registering data and specifying this method is also effective to trigger a process in the server to handle the posted data appropriately.

The passing of information to a process triggered to run on a server using either the GET or POST method is currently done according to an interface called the Common Gateway Interface (CGI). The receiving process is often written in a scripting language though this is not essential. Typically, the triggered server script is used for interfacing to a database to service a query included in a GET request. Another use, already referred to, is to append data associated with a POST request to a database.

Other important factors in the success of the WWW is the use of the HyperText Markup Language (HTML) for repre-

US 8,170,008 B2

9

senting the makeup of documents transferred over the WWW, and the availability of powerful graphical Web browsers, such as Netscape and Mosaic, for interpreting such documents in a client terminal to present them to a user. Basically, HTML is used to identify each part of a document, such as a title, or a graphic, and it is then up to the browser running in the client terminal to decide how to display each document part. However, HTML is more than this—it also enables a URI and a request method to be associated with any element of a document (such as a particular word or an image) so that when a user points to and clicks on that element, the resource identified by the URI is accessed according to the scheme (protocol) and request method specified. This arrangement provides a hyperlink from one document to another. Using such hyperlinks, a user at a client terminal can skip effortlessly from one document downloaded from a server on one side of the world, to another document located on a server on the other side of the world. Since a document created by one author may include a hyperlink to a document created by another, an extremely powerful document cross-referring system results with no central bureaucratic control.

Hyperlinks are not the only intelligence that can be built into an HTML document. Another powerful feature is the ability to fill in a downloaded "Form" document on screen and then activate a "commit" graphical button in order to have the entered information passed to a resource (such as a database) designed to collect such information. This is achieved by associating the POST request method with the "commit" button together with the URI of the database resource; activating the "commit" button results in the entered information being posted to the identified resource where it is appropriately handled.

Another powerful possibility is the association of program code (generally scripts to be interpreted) with particular documents elements, such as graphical buttons, this code being executed upon the button being activated. This opens up the possibility of users downloading program code from a resource and then running the code.

It will be appreciated by persons skilled in the art that HTML is only one of several currently available scripting languages delivering the functionality outlined above and it may be expected that any serious Web browser will have built-in support for multiple scripting languages. For example, Netscape 2.0 supports HTML 3.0, Java and LiveScript (the latter being Netscape proprietary scripting Language).

The importance of the role of the graphical Web browser itself should not be overlooked. As well as the ability to support multiple scripting languages, a Web browser should provide built-in support for standard media types, and the ability to load and execute programs in the client, amongst other features. These browsers may be viewed as operating systems for WWW interaction.

WWW and the Telephone Network

It is possible to provide a telephony service over the Internet between connected terminals by digitising voice input and sending it over the Internet in discrete packets for reassembly at the receiving terminal. This is an example of a communication service on the Internet. Conversely, it is possible to point to a variety of information services provided over the telephone system, such as the Minitel system widely available in France. However, these encroachments into each others traditional territories pose no real threat to either the Internet or the public telephone system.

Of more interest are areas of cooperative use of the Internet and the telephone system. In fact, one such area has existed for some considerable time and has been outlined above with

10

reference to FIG. 4, namely the use of a modem link over the PSTN from a user computer 34 to an Internet service provider 33 in order to obtain dialup IP access to the Internet. This cooperative use is of a very simple nature, namely the setting up of a bearer channel over the PSTN for subsequently generated Internet traffic; there is no true interaction between the Internet and the PSTN.

Another known example of the cooperative use of the Internet and PSTN is a recently launched service by which an Internet user with a sound card in his/her terminal computer can make a voice call to a standard telephone anywhere in the world. This is achieved by transferring digitised voice over the Internet to a service provider near the destination telephone; this service provider then connects into the local PSTN to access the desired phone and transfers across into the local PSTN the voice traffic received over the Internet. Voice input from the called telephone is handled in the reverse manner. Key to this service is the ability to identify the service provider local (in telephony charging terms) to the destination phone. This arrangement, whilst offering the prospect of competition for the telecom operators for long distance calls, is again a simple chaining together of the Internet and PSTN. It may, however, be noted that in this case it is necessary to provide at least a minimum of feedback to the Internet calling party on the progress of call set to the destination telephone over the PSTN local to that telephone; this feedback need only be in terms of whether or not the call has succeeded.

From the foregoing it can be seen that the current cooperative use of the Internet and telephone system is at a very simple level.

It is an object of the present invention to provide a method of accessing a service resource item over a communications network that facilitates the integration of the PSTN and the WWW.

SUMMARY OF THE INVENTION

According to the present invention, there is provided a method of accessing service resource items for use in respect of setting up bearer channels through a switched telecommunications system, the method including the steps of:

(a)—provisioning at least one server connected to a computer network with a plurality of service resource items that are thereafter locatable on said computer network by corresponding known URIs, said computer network being logically distinct from the telecommunications system, and said service resource items relating to setup control for bearer channels through said telecommunications system with each said service resource item being associated with a respective predetermined code, said predetermined codes being distinct from said URIs and identifying end-point entities for said bearer channels;

(b)—providing a mapping between each said predetermined code and the said known URI of the service resource item associated with that predetermined code; and

(c)—utilising a said predetermined code to access a corresponding said service resource item by using said mapping to determine the URI corresponding to that resource item and then using this URI to access the service resource item over said computer network.

In one embodiment, at least some of the URIs are derivable from their corresponding said predetermined codes by manipulation according to a function specified by said mapping. In another embodiment, at least some of the URIs are derivable from their corresponding said predetermined codes by look up in an association table associating said predetermined codes and URIs according to said mapping. This asso-

US 8,170,008 B2

11

ciation table can advantageously be held on at least one database server connected to the computer network, step (c) involving accessing the database server over the computer network to determine the URI corresponding to the said predetermined code. Preferably, the said at least one database server is provided by a DNS-type distributed database system in which the URIs are held in records associated with respective names, herein referred to as domain names, by which the records can be retrieved. In this case, step (c) involves translating said predetermined code into a corresponding domain name and using this domain name to retrieve the URI of the required service resource item from the DNS-type distributed database system.

More than one service resource item can be located at the same URI; in this case, the predetermined codes of these service resource items will include respective relative-resource-identifier values that can be used at the server holding the service resource items to identify the required resource item amongst the service resource items at the same URI.

The telecommunications system may be a telephone system with each said predetermined code being either the telephone number of the calling party or the telephone number of the called party (these numbers may either be the numbers of specific telephones, or personal numbers). In one preferred embodiment where at least some of said predetermined codes are called-party telephone numbers, the corresponding service resource items are the current telephone numbers of the called parties.

Generally as regards the nature of the service resources, these may be of the following type:

service logic intended to be executed by the corresponding server upon being accessed with the result of this execution being returned to the accessing entity;

downloadable service data which upon being accessed is intended to be downloaded to the accessing entity;

downloadable service logic which upon being accessed is intended to be downloaded to the accessing entity for execution thereby.

Preferably, where URIs are referred to in the foregoing, these URIs are URLs and/or URNs. Furthermore, the servers referred to are preferably HTTP servers.

It is to be understood that reference in the foregoing to the computer network being logically distinct from the telecommunications system is not to be taken to imply that there is physical separation of the two—indeed, there will frequently be joint use of the same physical infrastructure. Furthermore, not only may bearer channels set up in the telecommunications system share the same transmission medium as the computer network, but such a bearer channel may act as a pipe for traffic across the computer network. The intention of requiring the computer network to be logically distinct from the telecommunications system is to exclude computer networks that are dedicated to the management or monitoring of the bearer network and effectively form part of the telecommunications system itself.

Preferably, the computer network is generally accessible to users of the telecommunications system as this provides a number of benefits to users that will become apparent hereinafter. The phrase “generally accessible” should not be construed as meaning that all users of the telecommunications system have such access to the computer network or can get such access but, rather, it should be understood as meaning that a significant proportion of these users have or can obtain access to the computer network.

By way of example, in one preferred embodiment of the invention, the computer network generally accessible to users of the telecommunications system but logically distinct from

12

it, is the Internet and the telecommunications system is a public telephone system. In another embodiment, the telecommunications system is a private system including a PABX, and the computer network is a LAN.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will now be described, by way of non-limiting example, with reference to the accompanying diagrammatic drawings, in which:

FIG. 1 is a simplified diagram of a standard PSTN;

FIG. 2 is a simplified diagram of a known PSTN with IN service capability;

FIG. 3 is a diagram illustrating host domain name resolution by the DNS of the Internet;

FIG. 4 is a diagram illustrating the functioning of the World Wide Web;

FIG. 5 is a diagram illustrating the format of a standard URL;

FIG. 6 is a diagram of a first arrangement in which service resource items are held on HTTP servers accessible both by the service control subsystem of a PSTN and by Web users;

FIG. 7 is a diagram illustrating the processing of a service request by the SCP of FIG. 6;

FIG. 8 is a diagram illustrating the format of a resource code used by the FIG. 6 SCP when accessing a service resource item;

FIG. 9 is a diagram illustrating the process of accessing a service resource in the case where the service code does not include an RRI part;

FIG. 10 is a diagram illustrating the process of accessing a service resource in the case where the service code includes an RRI part;

FIG. 11 is a diagram illustrating the derivation of the URI of a service resource by parsing an input telephone number in accordance with the present invention;

FIG. 12A is a diagram depicting a name space (the “tel-name space”) constituted by the domain names derived by a parsing of a predetermined set of telephone numbers;

FIG. 12B is a diagram depicting the incorporation of the telname space without fragmentation into the DNS;

FIG. 12C is a diagram depicting the incorporation of the telname space in fragmented form into the DNS;

FIG. 13 is a diagram illustrating the overall operation of the FIG. 6 arrangement in providing a roaming number service in response to a telephone number being dialed at a standard phone;

FIG. 14 is a diagram illustrating the overall operation of the FIG. 6 arrangement when utilised by a Web user in setting up a call through a telephone interface integrated into the user's Web terminal;

FIG. 15 is a diagram illustrating the overall operation of an arrangement in which an interface is provided between the PSTN and the Internet for telephone traffic;

FIG. 16 is a diagram illustrating the overall operation of an arrangement in which a call setup gateway is provided between the Internet and the PSTN;

FIG. 17 is a diagram illustrating the overall operation of an arrangement in which a freephone service is implemented for Web users; and

FIG. 18 is a diagram similar to FIG. 6 illustrating the provision of a distributed processing environment for inter-connecting elements of the service control subsystem of the PSTN.

BEST MODE OF CARRYING OUT THE INVENTION

FIG. 6 illustrates an arrangement for the provision of services in a PSTN conventionally comprising an inter-ex-

US 8,170,008 B2

13

change network 13 (including trunks and switches at least some of which are SSPs 41 with associated IPs), an access network 11 connecting customer premise equipment (here shown as telephones 40) to the network 13, and a service control subsystem 42 including at least one SCP for providing services to the SSPs 41 upon request. It will be appreciated that the FIG. 6 representation of a PSTN is highly diagrammatic.

The SCP 43 may operate in a conventional manner responding to service requests from SSPs 41 to run specific service logic on particular data according to information contained in the service request, and to send back to the requesting SSP appropriate instructions for effecting call set up. A service request is generated by the SSP in response to predetermined trigger conditions being met at a trigger check point, there being one or more such check points in the course of handling a call (it may be noted that where the trigger conditions have been downloaded to the SSP from the SCP then it could be said that the SSP is responding to an information request by the SCP when contacting the SCP upon the trigger conditions being met—however, in the present specification, this initial communication from the SSP to the SCP will be referred to as a “service request”).

The SCP 43 is also provided with a network access interface 44 to the Internet 50 in order to make use of certain service resource items 49 (also referred to below simply as “service resources”) during the course of processing at least certain service requests from the SSPs 41. These service resources 49 are held as WWW pages on HTTP servers 51 (more particularly, on service resource databases 52 of these servers 51). The WWW pages containing these service resources are referred to below as “phone” pages. The servers 51 are connected to the Internet and the phone pages are read accessible using respective URLs or URNs (for convenience, the more general term URI will be used hereinafter to mean the Internet-resolvable indicator of the location of a phone page).

The service resources may be service logic or service data and may be used by an otherwise standard service logic program running on the SCP, by accessing the phone page of the required resource using the appropriate URI. In certain cases, the service resources 49 may provide substantially all of the service control and data associated with a particular service. In this case, the service logic program running in the SCP 43 is of skeleton form, being instantiated on receipt of a service request and then serving to initiate service resources access and to return the results of this access to the entity that made the service request. In fact, according to this approach, the SCP could be implemented simply as a platform for fetching and executing phone-page service logic and would not need to have the complex provisioning and management systems for such logic as is required by standard SCP platforms; SCPs could then become more ubiquitous, possibly being associated with every SSP.

FIG. 7 is a flow chart illustrating the progress of events in the case where the SCP 43 handles a service request by accessing a phone-page service resource. Upon receipt of a service request in an INAP message (step 100), SCP 43 decodes the TCAP/INAP message structure in standard manner (steps 101 and 102) well understood by persons skilled in the art. Next, SCP 43 instantiates a service logic program, SLP, to handle the request (step 103). This SLP is then responsible for looking up the URL of the required service resource as determined from information contained in the service request (steps 104, 105). For example, if the service request relates to a called-party service, then the required resource will be indicated by the dialed number and the latter will be

14

used to derive the URL of the resource. Once the URL of the desired service resource has been ascertained, a resource request (for example, in the form of an HTTP request message) is sent over the Internet to the corresponding server holding the desired service resource (step 106); a correlation ID is also passed with the resource request to enable a response from the latter to be linked with the appropriate SLP instance. A timer is also started (step 107).

If a response is received from the accessed resource before the expiration of a time-out period (tested in step 108), then the response, which is usually in the form of a destination number, is supplied to the appropriate SLP as identified using the correlation ID passed with the response (step 109). An INAP/TCAP response message is then prepared and sent to the entity that made the original service request (steps 110 and 111) after which the SLP instance is terminated (113).

If in step 108, a time-out occurs before a response is received, then a default response value (generally a default destination number) may be looked up in the customer record and put in an INAP/TCAP message and sent back to the requesting entity (steps 114 to 116). The SLP instance is then terminated (113).

Locating & Accessing Service Resources

The functionality associated with accessing a phone-page resource is schematically represented in FIG. 6 by resource access block 46. Block 46 includes URI determination block 47 for determining the URI of the phone page containing the desired resource on the basis of parameters passed to block 46. Using the URI returned by block 47, the resource access block 46 then accesses the phone page of the required service resource 49 over the Internet through interface 44.

Resource Codes—It is possible that more than one service resource is associated with a particular telephone number; in this case the resource access block 46 will need to know additional information (such as current point-in-call, pic) to enable the appropriate service resource to be identified. If the service resources associated with a number are located on different phone pages, then the additional information is also passed to the URI determination block 47 to enable it to return the URI of the appropriate phone page. It is also possible for all the service resources associated with a number to be located on the same phone page. In this case, the resource access block 46 uses the additional information to pass a resource-identifying parameter with its access request to the phone page concerned; it is then up to the functionality associated with the phone page to access the correct service resource.

Thus, each service resource can be considered as being identified by a respective resource code 54 (see FIG. 8) made up of a first part UI (“URI Identifier”) used to identify the URI at which the resource is located on the Internet, and a second part RRI (“Relative Resource Identifier”) used to identify the resource amongst plural resources at the same URI.

Resource Access—Where only one service resource 49 is located on a phone page 58 identified by a unique URI, then the resource code 54 simply comprises the UI, generally either a telephone number alone or a telephone number plus a pic parameter (see FIG. 9). In this case, accessing a resource simply involves mapping the whole resource code 54 into the corresponding URI (process 55) and then sending a request 57 to the corresponding phone page 58, this latter itself constituting the desired service resource 49. The result of accessing resource 49 is then returned in response message 59.

In contrast, where multiple service resources 49 are located on the same phone page 58 (FIG. 10), the resource code 54 comprises both a UI and RRI, the UI generally being a telephone number and the RRI a pic or other parameter for

US 8,170,008 B2

15

distinguishing between the co-located resources. In this case, accessing a resource involves mapping the UI part of the resource code 54 into the corresponding URI (process 55) and then sending a request 57 to the corresponding phone page (process 56), the request including the RRI of the resource code. The phone page 58 includes functionality 64 for accessing the required resource on the basis of the RRI in the request message. The result of accessing the required resource 49 is then returned in response message 59.

An alternative to the FIG. 10 method of accessing a service resource that is co-located with other resources on a phone page, would be to retrieve the whole page across the Internet simply using the URI derived from the UI part of the resource code, and then to extract the desired resource on the basis of the RRI.

URI Determination from Resource Code—The implementation of the URI determination block 47 that performs process 55 will next be considered. Block 47 may be implemented in a variety of ways, four of which are described below:

Direct Input

It would be possible, though not necessarily convenient, to arrange for the calling party to input directly the required URI. The calling party may thus input the host id component of the URI required (either in the form of a host domain name or host IP address) plus the path component of the URI. For example, in the case where the phone page of a called party is to be accessed, the calling party may input the URI of the called party and, indeed, this input may substitute for the normal input of a telephone number. A leading input string (for example "999") may be used to identify the input as an URI. As regards the input means, where a user only has a standard 12 key telephone, input of host domain names and other URI elements requiring alpha characters, will need to be done using one of the standard techniques for alpha input from a phonepad (such techniques are already used, for example, to enable a calling party to "spell" out the name of the called party). It would also be possible to provide users with a full alphanumeric keypad to facilitate URI input.

Computation

Service resource access over the Internet could be restricted to a set of dialed numbers from which it was possible to compute a corresponding URI; in this case, this computation would be the responsibility of block 47.

Association Table Lookup

Probably the simplest implementation for the block 47 is as an association table (either in memory or held on database disc store 48) associating a URI with the UI part of each resource code. A potential problem with this approach is that a service resource may be required for a called party number on the other side of the world which implies a rigorous update regime between PSTN operators worldwide in order to keep the association table up-to-date. (Note that the same implication is not necessarily applicable in respect of marking the called-party number as one required to trigger a service request, since the number may be arranged to be one of a group of numbers all triggering an appropriate service request, in a manner similar to 800 numbers).

DNS-Type Lookup

An alternative lookup solution is to use a hierarchically-structure-d distributed database system, similar to (or even part of) the Domain Name System (DNS) of the Internet, in order to resolve the UI part of a resource code to a corresponding URI. This approach, which will be described in more detail below, would typically involve databases maintained by each PSTN operator for its numbers with which URIs are associated. These databases would be accessible by all

16

PSTNs through a network such as the Internet with resolution requests being pointed to the appropriate database in a manner similar to the Domain Name System. In this case, the block 47 is constituted by an appropriate resolution program arranged to request UI resolution over the Internet through interface 44.

Before describing a DNS-type lookup implementation for the URI determination block 47, some further general comments are appropriate. Whatever method is used to determine the URI, certain simplifications are possible if limited constraints are placed on the URIs permitted. In particular, it is not necessary to determine all components of an URI in the following cases:

(i) A part of the URI path component can be made standard for all service resources, this standard part being simply added by the block 47 once the rest of the URI has been determined. For example, where a roaming number is to be looked up, it may by convention always be held in a file "roam" in a subdirectory "tel" of a subscriber's directory on a particular server. In this case the URI host component and the subscriber-unique part of the path component are first determined and then the remaining path part "/tel/roam" is added.

(ii) The URI path component can be arranged to be the same as a predetermined part of the resource code, the block 47 needing only to determine the host component and then add the path. For example, it may be agreed that the path must always end with the telephone number concerned, or sufficient of the terminating digits to have a high probability of uniqueness on the host machine. The path may also include standard components to be added by block 47.

(iii) Blocks of telephone numbers may have their corresponding service resources located on the same host server so that it is only necessary to use a part of the telephone number to determine the host component of the URI; in this case, the path component can conveniently include all or part of each telephone number. This situation implies a tight degree of control by the telephone operators and does not offer the telephone user the freedom to choose the host server on which user places their phone page.

Another general point worthy of note is that however the URI is determined, the host component of the URI may be provided either in the form of a host domain name or a host IP address. Where the host is identified by a domain name, then a further resolution of URI host name to IP address will subsequently be carried out in standard manner by interface 44 using the Domain Name System of the Internet. This further resolution can be avoided if the host identity is directly provided as an IP address.

Where a database lookup is used to provide the number to URI translation, this database may be independent of, or combined with, a customer database containing other customer-related information. Factors affecting this choice include, on the one hand, the possible desirability of having the number-to-UI translation information widely available, and on the other hand, the possible desirability of restricting access to other customer-related information.

DNS-Type URI Lookup

A DNS-type lookup implementation for the URI determination block 47 will now be described in some detail for the case where the UI part of the resource code is a telephone number and there are no constraints on the URI, thereby requiring both the full host and path components of the URI to be returned by the lookup. A key part of the overall process is the formation of the equivalent of a host domain name from the telephone number of interest; this domain-name equivalent is then resolved into a corresponding URI by a lookup

US 8,170,008 B2

17

mechanism which in the present example is identical to that employed by the DNS (indeed, the lookup mechanism may be incorporated into the DNS though it can also be independently implemented).

The nature of the DNS has already been described above with reference to FIG. 3 when the term "DNS-type" system was also introduced. For convenience in the following a DNS-type system organised to provide a telephone number to URI translation facility will be referred to as a "Duris" system (standing for "DNS-type URI Server" system).

The basic principles surrounding operation of a Duris system are:

every telephone number can be turned into a host domain name (the name space containing such host domain names for the telephone numbers of interest is referred to below as the "telname space"); and

for every host domain name in the host domain space there is a Registration Record held by the Duris system containing the corresponding URI.

Thus, an input telephone number forming, in the present case, the UI part of a resource code 54 (see FIG. 11), is first parsed to form a domain name (step 120) and then passed to the Duris system (illustrated in FIG. 11 as formed by the DNS itself) to retrieve the RR with the corresponding URI (step 121). Following on from the URI lookup, if the URI returned has its host component as a domain name, the DNS is next used to derive the host IP address (step 122); this step is, of course not needed if the host component is stored as an IP address in the RR. The URI is then used to make a resource request to the appropriate server, passing any RRI part of the resource code 54 (step 123).

There are a number of possibilities at the top level as to how a Duris system could be implemented:

(a) Independent of the DNS. In this option, the telname space constitutes the entire name space to be managed with the root of the telname space being the "." name space root (see FIG. 12A where the telname space is shown hatched). In this case, the Duris system is independent of the DNS itself. The Duris system could, of course, use the same basic infrastructure as the DNS (that is, the Internet) or an entirely separate network. Where the telname space comprises all the domain names corresponding to all public telephone numbers worldwide, parsing a full international telephone number would give a fully qualified domain name. Of course, the telname space could be a much smaller set of names such as those derived from internal extension numbers within a company having worldwide operations.

(b) Unfragmented Telname Space within the DNS. In this option, the telname space is a domain of the DNS name space and the Duris system is provided by the DNS itself. Thus, where the telname space comprises all domain names derived from public telephone numbers worldwide, the telname space could be placed within the domain of the ITU, in a special subdomain "tel", the root of the telname space then being "tel.itu.int." (see FIG. 12B where again, the hatched area represents the telname space). The responsibility for administering the domain "tel.itu.int." would then lie with the ITU. With this latter example, to form a fully qualified domain name from an input telephone number, after the number has been parsed to form the part of the domain name corresponding to the structuring within the telname space, the tail "tel.itu.int." is added. The fully qualified domain name is then applied to the DNS and the corresponding RR record, holding the required URI, is retrieved. As a further example, the telname space could be all name derived from internal extension numbers within Hewlett-Packard in which case the root

18

of the telname space would be "tel.hp.com." and Hewlett-Packard would be entirely responsible for managing this domain.

(c) Fragmented Telname Space within the DNS. In this option, the telname space is split between multiple domains of the DNS name space and the Duris system is provided by the DNS itself. Thus where the telname space comprises all domain names derived from public telephone numbers worldwide, the telname space could be split between respective "tel" subdomains of each country domain; thus, as illustrated in FIG. 12C, the part of the telname space corresponding to French telephone numbers would have a root of "tel.fr." and the part of the telname space corresponding to UK telephone numbers would have a root of "tel.uk.". The responsibility for administering each "tel" subdomain would then lie with each country. With this latter example, to form a fully qualified domain name from an input telephone number, the part of the telephone number following the country code is parsed to form the part of the domain name within a country tel subdomain and then a host domain name tail is added appropriate for the country concerned. Thus for a French telephone number, the "33" country code is stripped from the number prior to parsing and used to add a tail of "tel.fr.". The tail appropriate to each country can be stored in a local lookup table. As a further example, two commercial organisations (X company and Y company) with respective DNS domains of "xco.com." and "yco.com." may agree to operate a common Duris system with a telname space split between "tel.xco.com." and "tel.yco.com.". In this case, any Y company telephone number input from X company will be parsed to a fully qualified domain name terminating "tel.yco.com." and vice versa.

Consideration will next be given to the parsing of a telephone number into a domain name—in other words, where to insert the "." characters into the number to provide the structuring of a domain name. Of course, as already explained, telephone numbers are hierarchically structured according to each country's numbering plan. Thus one approach would be to follow this numbering plan structuring in dividing up a telephone number to form a domain name. By way of example, the telephone number "441447456987" which is a UK number (country code "44") with a four digit area code ("1447") and six digit local number ("456987") could be divided to form a domain name of 456987.1447.44 (note that the reversal of label order occasioned by the fact that the DNS labels are arranged least significant first). If the telname space is a subdomain of the DNS with a placement as illustrated in FIG. 12B, the fully qualified domain name derived from the telephone number would be:

456987.1447.44.tel.itu.int.

There are however, difficulties inherent with trying to match the numbering plan hierarchy when parsing a telephone number into a host name. Firstly, in order to parse an international number correctly, it would be necessary for each entity tasked with this operation to know the structuring of each country's numbering plan and where, as in the UK, area codes may be of differing length the required knowledge may need to take the form of a lookup table. Whilst this is not a complicated computational task, it is a major administrative nuisance as it means that each country will need to inform all others about its numbering plan and any updates. The second problem is that a six or seven digit local number is a very large domain; it would be preferable to create subdomains for performance and scaling reasons but there is no obvious way of doing this.

These problems can be overcome by giving up the restriction that the parsing of telephone number into a domain name should match the structuring of national numbering plans. In

US 8,170,008 B2

19

fact, there is no strong reason to follow such a scheme as DNS servers know nothing about the meaning of the name space. It is therefore possible to parse telephone numbers using a deterministic algorithm taking, for example, 4 digits at a time to limit the size of each subdomain and making it possible to 'insert the dots' without knowing the numbering plan concerned. So long as the DNS domains and zones served by the DNS servers are created correctly it will all work.

For international numbers it would still seem appropriate to separate off the country codes and so a hybrid parsing scheme would be to parse the initial part of a dialed number according to known country codes and thereafter use a deterministic scheme (for example 3, 7 or 4, 6 or 3,3,4) to separate the digits. Of course, if a fragmented telname space is being used as illustrated in Figure UC then the country code is used to look up the host name tail and it is only the national part of the number which would be parsed.

Finally, as regards the details of how a DNS server can be set up to hold RR records with URIs, reference can be made, for example, to "DNS and BIND", Paul Albitz and Cricket Liu, O'Reilly & Associates, 1992 which describes how to set up a DNS server using the Unix BIND implementation. The type of the RR records is, for example, text.

It should be noted that DNS labels should not in theory start with a digit. If this convention is retained, then it is of course a trivial exercise when parsing a telephone number to insert a standard character as the first character of each label. Thus, a 4 digit label of 2826 would become "t2826" where "t" is used as the standard starting character.

It will be appreciated that as with domain names, where an input telephone number is not the full number (for example, a local call does not require any international or area code prefix), it would be parsed into a domain name in the local domain.

The foregoing discussion of Duris system implementation, has been in terms of translating a telephone number into an URI where the telephone number forms the full UI of a resource code and the Duris system returns a full URI. It will be appreciated that the described Duris implementation can be readily adapted to accommodate the various modification discussed above regarding the form of the UI and what parts of the URI need to be looked up. For example, where there are a number of different service resources associated with a subscriber each in its own file and the required source is identified by a pic part of the resource code, then the input telephone number will be used to look up, not the full URI, but the host component and that part of the path component up to the relevant subdirectory, the pic part of the UI then being appended to identify the required resource file.

For small local Duris implementations, it may be possible to have a single server; the implementation should still, however, be considered as of a DNS type provided the other relevant features are present.

Nature of Service Resources

Turning now to a consideration of the service resources, how these service resources can be provisioned onto the servers will be described more fully below but, by way of present example, the service resource or resources associated with a particular PSTN user (individual or organisation, whether a calling or called party) can be placed on a server over the Internet from a user terminal in one or more WWW pages.

Consider the simple case where the service resource is a service data item such as a telephone number (for example, an alternative number to be tried if the user's telephone corresponding to the number dialed by a calling party is busy). This diversion number could be made the sole service resource of

20

a phone page of the user. The phone page URI could be a URL with scheme set to HTTP in which case the GET method could be used to retrieve the diversion number. Such an arrangement is suitable if the phone page is only to be used for functional retrieval of the diversion number. However, if the diversion number is to be visually presented at a user terminal, then it may be desirable to accompany the number with explanatory material (this will often not be necessary as the diversion number can be arranged to be returned into an existing displayed page that already provides context information). However, where the phone page does include explanatory material as well as the diversion number, an entity only wishing to make functional use of the phone page, could be arranged to retrieve the phone page and then extract the diversion number (this would, of course, require a standard way of identifying the information to be extracted from the phone page).

An alternative and preferred arrangement for providing for both viewing and functional access to a resource requiring explanatory material for viewing, is to use an object-oriented approach to resource design. In this case, the resource object would have two different access methods associated with it, one for purely functional use of the resource and the other enabling viewing of associated explanatory material. It would then be up to the accessing entity to access the resource object using the appropriate object method.

Yet another arrangement for providing for both viewing and functional use of the diversion number, would be to provide separate resources appropriately configured for each use, each resource having its own resource code (generally, both such resources would be placed on the same phone page and in this case the UI part of each resource code would be the same).

Retrieval of a phone page for use by a human user will generally not be as time critical as retrieval for operational use by a PSTN. Thus, while for human use the scheme specified in the URL of a service resource could be HTTP, it may be advantageous for operational use to define a special "phone" scheme (access protocol) which would result in the server using an optimised access routine to access the required resource (diversion number, in the current example) and respond to the accessing entity in the minimum possible time.

Besides data items, other possible types of service resource include service logic for execution in place (at the server) with the result of this execution being returned to the entity accessing the resource; service logic downloadable from the server to the accessing entity for execution at that entity; and a logging resource for logging information passed to it by the accessing entity (or simply for logging the fact that it has been accessed). It will be appreciated that the logging resource is really just a particular case of service logic executable in place.

By way of example, a service resource constituted by execute-in-place service logic can be arranged to implement time-of-day routing, the result of executing the service logic being the telephone number to which a call should be routed taking account of the time of day at the called party's location. An example of a service resource constituted by downloadable service logic is service logic for controlling calling-party option interrogation using the facilities provided by an IP. As regards the logging resource, this can be used for recording the number of calls placed to a particular number.

Where each resource has its own phone page and the resource is present only in its unembellished functional form, then the HTTP scheme can be employed for access using the GET method for both the downloadable service logic and the execution-in-place service logic, and the POST method for

US 8,170,008 B2

21

the logging resource. If it is desired to provide an explanatory material with each service resource, then any of the solutions discussed above in relation to data items, can be used.

Where more than one service resource is to be associated with a number, then each such resource can be placed on a respective phone page with its own URI. However, the preferred approach is to place all such service resources on the same page and use the RRI part of the corresponding resource codes to enable access to the appropriate resource. The accessed resource is then treated according to its form (executed if execute-in-place service logic, returned if downloadable service data or logic).

Thus if both a diversion-number service-data resource and a time-of-day execution-in-place service-logic resource are placed on the same phone page, the diversion-number resource code might have an RRI of "1" whilst the time-of-day resource code might have an RRI value of "2".

Where calling/called party options are to be included in a service resource for presentation to such party, then as already indicated, this can conveniently be done by constituting the service resource as downloadable service logic with the chosen option possibly initiating request for a follow-up service resource.

It will be appreciated that a service resource will often be of a complex type, combining service data and/or downloadable service logic and/or execute in place service logic. A particularly powerful combination is the combination of the two types of service logic where the downloadable service logic is designed to interact with execute-in-place service logic; using this arrangement, the user can be presented with complex client-server type applications.

Example Usage of Service Resource

FIG. 13 illustrates the operation of a service making use of a resource on a server 51. This service is equivalent to a "personal number" service by which a user can be accessed through a single, unchanging number even when moving between telephones having different real numbers. To achieve this, the user requiring this service (user B in the current example) is allotted a unique personal number (here referred to as the "Webtel" number of B) from a set of numbers all of which have the same leading number string to enable an SSP to readily identify a dialed number as a Webtel number. User B has a service resource 49 on a dedicated phone page on HTTP server 51, this phone page being located at a URL here identified as "URL (B phone page)". B's phone page when accessed returns the current roaming number ("B-telNb") where B can be reached. In the simplest case, B's phone page is just a single number that can be modified by B (for example, from a terminal 53) as B moves to a different phone. More likely is that B's phone page is an execute-in-place service logic providing time of day routing.

In the present example, the association between B's Webtel number and the URL of B's phone page is stored in an association table accessible to SCP 43.

Upon a user A seeking to contact user B by dialing the Webtel number of B, the telephone 40 being used by A passes a call set up request to SSP 41 (note that in FIG. 13 the bearer paths through the telephony network are shown by the thicker lines 60, the other heavy lines indicating signalling flows). SSP 41 detects the dialed number as a Webtel number and sends a service request to SCP 43 together with B's Webtel number. SCP 43 on receiving this service request initiates a service logic program for controlling translation of B's Webtel number into a current roaming number for B; in fact, in the present case, this program simply requests the resource access block 46 to access the service resource identified by B's Webtel number, (that is, B's phone page 49) and return the

22

result of this access. To this end, block 46 first translates B's Webtel number into the URL of B's phone page and then uses this URL to access B's phone page over the Internet (for example, using the phone scheme already referred to with a method corresponding to the HTTP GET method). This results in B's current roaming number B-telNb being passed back to block 46 and in due course this number is returned to the SSP 41 which then initiates completion of call set up to the telephone 40 corresponding to B-telNb.

The FIG. 13 example related to a called-party service; it will, of course, be appreciated that the principle of accessing service resources over the Internet can be applied to all types of services, including both calling-party and called-party services and hybrids. Thus, standard 800 number services can be implemented with the dialed 800 number resulting in access to a phone page resource constituted by execute-in-place service logic that returns the most appropriate number for controlling onward call routing.

It will be appreciated that although in the FIG. 13 example the service request from the SSP was triggered by a leading number string of a dialed number, a service request may be triggered by a variety of triggers including calling-party number, called-party number, or some other user input, such triggers being possibly qualified by call setup progress (for example, called-party number qualified by a busy status or by ringing for more than a certain time).

With respect to the logging service resource mentioned above, one possible application for such a resource is in telephone voting. In this case, dialing the voting number causes the SSP picking up the call to pass a service request to SCP 43 which then contacts the appropriate logging resource over the Internet to register a vote after which the call is terminated. To minimise bottlenecks, a logging resource could be provided at a different URL for each SCP, it being a simple matter to collect and collate voting from all these logging resources over the Internet. If an SCP with Internet access is provided at every SSP, then the risk of congestion is greatly reduced.

As already noted, a user's phone page may hold multiple service resources in which case the access request from the accessing SCP needs to contain an appropriate RRI identifying the required resource.

In the event that an SCP is to provide both a traditional IN service to some users and an equivalent service using an Internet-accessed service resource to other users, then a lookup table may need to be provided in the SCP to ensure that a service request is appropriately handled; such a lookup table can conveniently be combined with the customer record database.

Once a user, such as user B, has set up one or more phone pages specifying his desired service resources (particularly service logic defining personalised services), it is clearly logical for user B to want any PSTN operator he cares to use, to access and utilise such service resources. This is possible if the Webtel-to-URI databases are available to all operators. Thus multiple operators could be set to access B's phone page or pages. If an operator declines to use B's phone pages, B can obviously choose not to use that operator (at least where that operator provides a long haul carrier service subject to user selection). The possibility therefore arises that service provision will cease to command a premium from operators, but that the provision of phone-page utilisation by an operator will become a necessary basic feature of PSTN operation.

Provisioning and Updating Service Resources

Consideration will next be given as to how the service resources 49 are provisioned to the servers 51 and subsequently updated.

US 8,170,008 B2

23

So far as provisioning is concerned, two basic actions are required: firstly, the service resource must be placed on a server 51 and, secondly, the URI of the service resource must be notified to the PSTN operator along with the trigger conditions (number plus any other condition such as point in call) calling for access to the resource; if multiple resources are provided at the same URI, then the RRI values needed to retrieve the appropriate resource for a particular trigger condition, must also be notified. This notification process will be referred to hereinafter as registering the service resource with the PSTN operator; registration is, of course, necessary to enable the association tables used by SCP 43 to be set up and for trigger conditions to be set in SSPs 43. For certain services, such as that described above with reference to FIG. 13, it is not the user that supplies the triggering number (the Webtel number in the FIG. 13 example); instead, the PSTN operator allocates an appropriate number to the user as part of the registration process.

As to the process of placing a service resource on a server 51, how this is carried out will depend on the attitude of the PSTN operator to the possible effects of such service resources on operation of the PSTN. Where the service resource simply returns a data item to an accessing entity, then an operator may not be too concerned about possible errors (accidental or deliberate) in implementing the service resource. However, the operator will probably be much more concerned about the proper operation of any service logic that may be returned by a resource; indeed, an operator may not permit such a service resource.

Assuming for the moment that an operator has no concerns about the nature or implementation of service resources, then how a resource is placed on a server 51 will largely depend on the nature of the server concerned. For example, if a user has a computer with network access to the Internet and this computer is used as server 51, then the user can simply load a desired resource onto the server as a WWW phone page for external access. A similar situation arises if the server is an organisation server to which the server has access over an internal LAN. In both these latter cases, loading the resource as a WWW phone page does not itself require Internet access. However, if the server 51 is one run by an external Internet service provider, then a user can arrange to download the required service resource into the user's allocated Web site space on the server; this may or may not involve Internet access. One special case of this latter scenario is where the PSTN operator provides a special server for user phone pages containing service resources.

Except where a user's own computer acts as server 51, placing a service resource on a server will generally involve clearing one or more levels of password protection.

As regards the origin of the service resource loaded by a user onto server 51, this may be generated by the user or, particularly where the resource includes service logic, may be provided by a third party (including the PSTN operator).

If the PSTN operator wishes to have control over the service resources 49 to avoid any adverse effects on operation of the PSTN, two approaches are possible. Firstly, the operator could require that every resource (or, possibly, a particular subset) had to be subject to a verification process before use, appropriate measures then being taken to avoid subsequent alteration of the resource by the user (except, possibly, for particular data items); in this respect, the operator could require that the resource be placed on a server under the operator's control and to which the user had no write access (except possibly for altering particular data items, as indicated above). A second, more attractive, approach to minimising adverse effects by the service resources 49, is for the

24

operator to provide standard service resources to which a user could add the user's own data (and possibly make limited functional selections in case where the resource included service logic); the customised resource would then be loaded onto a server 51 controlled by the operator. This process can be conveniently implemented for a particular resource using an HTML "form" which a user could download over the WWW from the operator-controlled server. After completing the form and activating a commit graphical button of the form, the entered information would be posted back to the server where the information would be used to produce a customised service resource thereafter placed on the server for access over the Internet. An advantage of this approach is that registration of the service resource with the operator is simultaneously effected. (It may be noted that if registration needs to be done as a separate act from having a service resource loaded on a server, then using an HTML form is a very convenient way to implement the registration process).

From the foregoing it can be seen that whilst the provisioning process does not necessarily require information to be passed over the Internet, in many cases this will be the best solution, particularly if an HTML form exchanged over the WWW can be used to produce a customised service resource. It should be noted that producing a customised service resource using an HTML form is not limited to cases where the PSTN operator controls the server.

As regards updating service resources, there is likely to be a need to update certain data items on a fairly frequent basis (for example, roaming number). Where the PSTN operator does not place any controls on the service resources 49, then update is a relatively simple matter, only requiring write access to the server concerned (as already indicated, this will generally involve one or more levels of password protection). However, where the PSTN operator exercises control over the service resources, for example by only permitting customisations of standard service resources, such customised resources being loaded on servers controlled by the operator), then write access to the service resource may be tightly controlled. Again, an HTML form may conveniently be used as the medium for modifying a data item in such cases; to the operator, this has the benefit of limiting the modifications possible whilst to the user, a form interface should provide a simple route to resource modification.

For more complex updates, it may be necessary to go through a process similar to that required for initial provisioning.

Particularly where the service resources are held on a server 51 controlled by the PSTN operator, resource update will generally involve communication over the Internet.

Web User Interaction

Consideration will next be given to other possible uses of the service resources held in phone pages on the servers 51. For example, if user B's phone page contains a diversion number, then provided this phone page is read-accessible over the Internet from user A's terminal 53, user A can use a graphical Web browser running on terminal 53 to view B's phone page and discover B's diversion number. As earlier discussed, the diversion number may be passed to user A for display in an existing visual context giving meaning to the number, or may be passed to user A with accompanying explanatory text.

A more useful example is a current roaming number service for user B. Suppose B's phone page 49 on server 51 (see FIG. 14) is operative when accessed to return a current roaming number where B can be reached. Further suppose that user B has a Web site with several Web pages written in HTML and each page contains a graphical phone button which when

US 8,170,008 B2

25

activated uses the GET method to access B's phone page by its URL. Now if user A whilst browsing (arrow 66) B's Web site over the WWW from user A's terminal 53, decides that he would like to call user B to discuss some item of interest, user A simply activates the phone button 65 on the currently viewed page of B. This causes B's phone page to be accessed using the HTTP request "GET URL (B Phone Page)"—see arrow 67.

B's current number to be called is then determined and passed to user A's terminal 53 (see arrow 68) where it is displayed. An explanatory text concerning the number will generally also be displayed; for example the text "Please call me at the following number:" could be displayed, this text being provided either by the HTML script associated with the phone button, or from the phone page when returning the current number. In fact, it would probably be more helpful to provide user A, not only with the current number for reaching user B, but also with all numbers where B could be reached together with the times when B was most likely to be at each number. Since this extra information is likely to be subject to frequent change, the only sensible way to provide the information is from the phone page. Thus, B's phone page not only provides the current number for reaching B, but also a text that includes numbers and times subject to change; scripting B's phone page is, of course, done in a way that ensures that variable data need only be altered in one place.

In a further example, B's phone page might include downloadable service logic for execution at user A's terminal. This is useful where choices are to be presented to a user, each choice producing a follow-up action such as fetching a further phone page. For example, the first-accessed phone page may be a family phone page giving the general telephone number for a family but also giving the user the possibility of selecting further phone information on each family member, such as a time-of-day dependent number; in this case, each family member has their own follow-up phone page.

In the above scenarios, user A has been presented with a number to call over the PSTN. User A can now pick up his standard telephone and dial the number given. In fact, a complication arises if A only has Internet access via a SLIP/PPP connection over an ordinary, non-ISDN, PSTN line since, in this case, A's telephone line is already tied up with making Internet access when gateway 90 seeks to set up a call to A's telephone; with an ISDN connection, as two channels are available, this problem does not arise. One way of overcoming this problem would be to have user A's terminal 53, after obtaining the number to call from B's phone page, automatically suspend its Internet session by storing any required state information (for example, current WWW URL being accessed) and then terminate its SLIP/PPP connection to thereby free up the telephone line. A can then telephone B. At the end of this call, A can resume the suspended Internet session, using the stored state information to return to the point where A left off to call B. An alternative approach is to operate a suitable multiplexing modulation scheme on the telephone line to A allowing voice and data to be simultaneously carried. A number of such schemes already exist. The PSTN would then need to separate the combined data and voice rams coming from A at some point and pass each to its appropriate destination (the Internet data being forwarded to the ISP providing the SLIP/PPP connection for user A and the voice stream being passed to B); of course, data and voice traffic in the reverse direction would also need combining at some point for sending over the last leg to A's terminal.

Rather than A manually dialing B using a standard telephone, another possibility is that user A's terminal is provided with functionality enabling A to make a call over the PSTN

26

from his terminal; this functionality generally comprises a hardware interface 70 (FIG. 14) to a telephone line and phone driver software 71 for driving the interface 70 in response to input from application software such as the Web browser 73. A could call up his phone software and enter the required number or, preferably, A need only "select" on screen the number returned from B's phone page and then pass it into A's phone software. Indeed, provided user B knew the software interface to the software 71 providing dialing functionality on A's terminal, it would be possible for B's phone page to return to A's terminal program code for automatically dialing B's number upon A confirming that he wishes to proceed with call placement. As an alternative to placing a voice call, if A's terminal is equipped with a suitable modem and controlling software, A could, instead, elect to send a fax or data to B through the PSTN either to B's ordinary number or to one specified in B's phone page as the number to be used for such transmissions. Of course, placing a call from A's terminal over the PSTN may be subject to the problem already discussed of conflict for use of the telephone line where this is not an ISDN line and A gains Internet access via a SLIP/PPP connection.

However the call is placed, if B's telephone corresponding to the number tried by A is busy, a number of possibilities exist. Thus if B has a phone page that specifies a diversion number, and B has registered this service resource with the PSTN, then the diversion number should be automatically tried by the PSTN. However, if the diversion number resource has not been registered with the PSTN a busy signal will be returned to A. Where A has placed the call through a standard telephone, A must now decide how to proceed and A may elect either to give up or to refer again to B's phone page to look up the diversion number and redial using this number. If A placed the original call using his terminal 53 then the latter can be programmed to detect the return of a busy signal and then automatically look up B's diversion number and redial using this number. This functionality can be included in service logic downloaded from B's phone page and run on A's terminal.

If A had to terminate his Internet session in order to free up the telephone line for voice use, then referring back to B's phone page requires a new Internet session to be started (in fact, this inconvenience could be avoided if B's diversion number were passed to A's terminal at the time the original number to be dialed for B was supplied).

The service resource accessed on B's phone page upon B's telephone being busy may, of course, be more complex than just a diversion number. In particular, user A may be presented with a range of options including, for example, B's fax or voice mailbox number, the selection of an option potentially initiating the running of appropriate accessing software. Another possible option would be for A to leave B a call back message using a form downloaded from B's phone page upon this option being chosen; the completed form would be posted back to server 51 and logged for B to check in due course.

Of course, it may arise that user A wishes to access B's phone page to find out, for example, B's current roaming number, but user A does not know the URI of B's Web site and only has B's Webtel number. A could just call B through the PSTN in which case the translation of B's Webtel number to roaming number would be automatically effected (assuming B is still registered for this service); however, A may not wish to call B straight away, but just note his current roaming number. In order to solve A's problem, the Webtel-to-URI association tables previously described are preferably made accessible on the Internet at a known address (for example, at a known Web site). All that A need now do is to access this

US 8,170,008 B2

27

Web site passing B's Webtel number; B's phone page URI will then be returned to A who can then use it to access B's phone page. This process can, of course be made automatic from the point when A sends B's Webtel number to the association-table Web site.

Internet/PSTN Call Interface

In the FIG. 14 scenario, A's access to the PSTN was through a standard telephone interface even though the actual form of A's telephone differed from standard by being integrated into A's computer terminal 53. FIG. 15 illustrates a situation where A, after being supplied with B's current roaming number as in the FIG. 14 case, calls B via a route that starts out over the Internet and then passes through a user network interface 80 into the PSTN. Interface 80 is arranged to convert between ISDN-type telephone signalling on the PSTN and corresponding signalling indications carried across the Internet in IP packets; in addition, interface 80 transfers voice data from IP packets onto trunk 60 and vice versa.

Thus, upon A initiating a call to B, Internet phone software 81 in A's terminal sends call initiation signalling over the Internet to interface 80, the address of which is already known to A's terminal. At interface 80, the signalling is converted into ISDN-type signalling and passed to SSP 41. Call set up then proceeds in the normal way and return signalling is transferred back through interface 80, over the Internet, to the software 81 in A's terminal. This software passes call setup progress information to the WWW browser 73 for display to A. Upon the call becoming established, A can talk to B through his telephone and A's voice input is first digitised in phone hardware interface 83 and then inserted into IP packets by software 81 to traverse the Internet to interface 80 (see arrow 84); voice traffic from B follows the reverse path.

IN services can be provided to this call by SCP in response to a service request from an SSP 41. Thus, if B's phone is busy, and B is registered for call diversion, SCP 43 on receiving a service request will access B's appropriate phone page for call diversion and retrieve the diversion number. If SSP 41 is not set to initiate a service request on B's telephone being busy, the busy indication is returned to A's terminal where it can be handled in the manner already described with reference to FIG. 14.

In fact, interface 80 can be provided with functionality similar to an SSP to set trigger conditions and generate a service request to SCP 43 on these conditions being satisfied.

Third-Party Call Setup Gateway

FIG. 16 illustrates a further arrangement by which A can call B after receiving B's current roaming number. In this case, a third-party call set-up gateway 90 is provided that interfaces both with the Internet 50 and with an SSP 41. Conveniently, gateway 90 can be co-located with SCP 43 (though this is not essential). Gateway 90 has the capability of commanding SSP 41 to set up a call between specified telephones.

Thus, upon A wishing to call B, a third-party call setup request is sent from A's terminal over the Internet to the gateway 90 (see arrow 91). This setup request includes A's telephone number and B's current roaming number. Gateway 90 first attempts to setup the call to A's telephone (which should generally succeed) and thereafter to set up the call to B's identified telephone. Once the call is setup, A and B communicate in standard manner across the PSTN.

If B's phone had been busy, then any of the previously described scenarios may ensue.

Gateway 90 can also be arranged to make service requests to SCP 43 upon predetermined trigger conditions being satisfied. Thus, gateway 90 might be set to pick up the busy condition on B's telephone and initiate a service request to

28

SCP 43 for a diversion number. However passing the busy indication back to A's terminal via gateway 90 is preferred because of the flexibility it gives A regarding further action.

As already generally discussed in relation to FIG. 14, a complication arises if A only has Internet access via a SLIP/PPP connection over an ordinary, non-ISDN, PSTN line since, in this case, A's telephone line is already tied up with making Internet access when gateway 90 seeks to set up a call to A's telephone. The solutions discussed in respect of FIG. 14 (termination of Internet session; multiplexing voice and Internet data on same telephone line) can also be used here. An alternative approach both for FIG. 14 and for FIG. 16 scenarios is possible if user A's terminal can handle a voice call as digitised voice passed over the Internet. In this case, the voice call can be placed through an interface 80 of the FIG. 15 form, and the voice traffic and the Internet communication with the B's phone page and/or gateway 90 are both carried in Internet packets passed over the SLIP/PPP connection to/from A's terminal 53 but as logically distinct flows passed to separate applications running on terminal 53.

It may be noted that the third party call setup request made by A's terminal to gateway 90 could equally have been made by service logic held in B's phone page and executed by server 51 (such an arrangement would, of course, require A's telephone number to be passed to B's phone-page service logic and this could be arranged to occur either automatically or through a form presented to user A at terminal A and then posted back to server 51).

It may also be noted that the interface 80 of FIG. 15 and the gateway 90 of FIG. 516 provide examples of service requests being passed to the service control subsystem by entities other than SSPs 41.

WWW-Based "FreePhone" (800 Number) Services

It is possible to implement a "FreePhone" or "800 number" type of service using a combination of the WWW and the PSTN. As will be seen from the following description of such a service with reference to FIG. 17, a WWW/PSTN implementation does not necessarily rely either on transferring call charges from the calling to called party or on the use of a special "800" number, two characteristics of standard "Free-phone" schemes. The WWW/PSTN implementations do, however, possess the more general characteristic of placing an enquiring party and the party to whom the enquiry is directed, in telephone contact at the expense of the latter party.

In the FIG. 17 arrangement, a user D such as a large department store has a website on a server 51; for the sake of simplicity, it will be assumed that the server is under the control of user D who has direct computer access to the server over line 125. D's Website may, for example, contain many catalogue-like Web pages illustrating goods offered for sale by D. In addition, D has a freephone page 124 for handling enquiries placed on a freephone basis; the URL of this page is associated with a "freephone" graphical button 122 placed on each of the Website catalogue pages.

Suppose user A at terminal 53 is browsing D's Website, looking at the catalogue pages (arrow 121). If A sees an item of interest and wishes to make an enquiry to D about this item, then A can activate at terminal 53 the graphical freephone button 122 associated with the catalogue page concerned. This activation causes code embedded in the catalogue page currently loaded in A's terminal to prompt the user to enter their telephone number and, optionally, their name, after which an HTTP request is sent to D's freephone page using the POST method and enclosing the entered data (arrow 123). D's freephone page on receiving this request executes service logic to enter a new enquiry (including A's name and tele-

US 8,170,008 B2

29

phone number) in an enquiry queue 127 maintained in an enquiry control system 126. In the present example, the enquiry control system is connected to the server 51 via line 125, externally of the Internet; however, it would also be possible to have server 51 communicate with the enquiry control system through the Internet and, indeed, this may be the most practical arrangement where D's Website is on an ISP server rather than on a server controlled by D. In fact, the code run in A's terminal upon activation of the freephone graphical button 122 could be arranged to directly forward the enquiry request to the enquiry control system over the Internet rather than passing it back through the server 51.

The enquiry control system 126 manages enquiries passed to it to ensure that they are dealt with in an ordered manner. The system 126 on receiving a new enquiry preferably estimates approximately how long it will be before the enquiry is dealt with, this estimation being based on the number of currently queued enquiries and the average time taken to handle an enquiry. This estimation of waiting time is passed back via server 51 to user A in the response to the POST request message.

The enquiry control system 126 looks after the distribution of enquiries to a number of agents each of which is equipped with a telephone 40 and a display 129. A's enquiry will be dealt with as soon as it reaches the head of the queue 127 and there is an agent detected as available to handle the enquiry (thus, for example, the system may be arranged to detect when an agent's telephone goes on hook). When these conditions are met, a distribution and setup control unit 128 takes A's enquiry and displays A's name and telephone number on the display 129 of the available agent (for clarity, herein referenced as agent D'); if user D keeps a database on D's past customers or credit rating data, then unit 128 will also look for and display any such further information known about A. At the same time, unit 128 makes a third-party call setup request (arrow 130) over the Internet to gateway 90 asking for a call to be set up between the telephone of the available agent D' and the telephone of user A, both telephones being identified by their respective numbers. If both D' and A pick up the call, the enquiry then proceeds, the cost of the call being paid for by D as it is D that originated the call over the PSTN. If, for whatever reason, the call remains incomplete (for example, unanswered by A) for a predetermined timeout period, then unit 128 can be arranged to automatically pass on to the next enquiry at the head of the queue 127.

It would, of course, be possible to dispense with having the unit 128 request call setup through gateway 90 and either have the agent D' dial A's number manually or have unit 126 initiate auto-dialing for D' telephone (agent D' having, for example, a computer-integrated telephone similar to that of A's in FIG. 14). The advantage of these approaches is that the existing PSTN could be used without adaption and without any service installation, in implementing the WWW-based freephone service.

As discussed in relation to FIGS. 11 and 13, a complication arises in placing a call to A if A only has Internet access via a SLIP/PPP connection over an ordinary, non-ISDN, PSTN line since, in this case, A's telephone line is already tied up with making Internet access when user D tries to set up a call to A's telephone. The solutions discussed in respect of FIGS. 11 and 13 can also be used here (termination of Internet session; multiplexing voice and Internet data on same telephone line; and placing the call over the Internet to A's terminal). With respect to the solution based on termination of the Internet session, such termination could be delayed until A's enquiry was about to be dealt with; however, to do this, it would be necessary to provide feedback from the control

30

system 126 over the Internet to A's terminal 53 and to associate this feedback with code for bringing about Internet-session termination. One way to achieve this would be to have the response message sent by server 51 in reply to the original POST request message from A, include a correlation code; any subsequent feedback from system 126 passed to A would also include this code (server A having also passed the code to control system 126) thereby allowing A's terminal to correctly identify this feedback. In fact, the same mechanism could be used to provide user A with updates on how much longer user A is likely to be waiting to be called back, this mechanism being usable independently of whether or not there was a conflict problem for use of A's telephone line.

Where user A only has a telephone 40 and no terminal 53, it is still possible to utilise the basic structure of FIG. 17 to provide a freephone service for user A without resorting to the complexity of call charge transfer. More particularly, A would dial a special number for user D's freephone service (typically an 800 number), and the SSP 41 would recognise this special number in standard manner and make a service request to SCP 43 including both this special number and A's number. SCP 43 would then ascertain D's freephone-page URL by doing a number-to-URL translation and access D's freephone page using a POST-method HTTP request similar to request 123. Once this request had been registered as an enquiry by D's freephone page 124, the latter could send a response to SCP 43 asking it to play an announcement such as "Your freephone enquiry has been registered; please hang up and you will be contacted shortly". This announcement could be played to A by an IP in standard manner. A would then hang up and be ready to receive a call from D.

A significant advantage of the above freephone schemes using WWW, is that user D is not running up charges for use of the PSTN during periods when an enquiry is enqueued, waiting to be handled.

Variants

Many variants are, of course, possible to the above-described arrangements and a number of these variants are described below.

Distributed Processing Environment. As is illustrated in FIG. 18, the SCP 43 may access the HTTP servers 51 through a distributed processing environment, DPE 98, at least logically separate from the Internet. Preferably in this case the servers 51 are controlled by PSTN operators and are thus restricted in number.

Service Resources on DNS-Type Servers. In the foregoing examples, the service resource items have been placed on servers 51 connected to the Internet and a desired service resource has then been accessed over the Internet by the service control subsystem of the PSTN, and/or by Internet users, through the use of an URI derived from a resource code that identifies the desired service resource item. In a preferred arrangement for deriving the URI from a resource code in the form of a telephone number, all or part of the telephone number concerned was parsed into domain name form and then resolved into an URI using a DNS-type distributed database system that, indeed, could be integrated into the DNS itself (see FIGS. 11 and 12, and related description). In fact, it would be possible to place service resource items directly in Registration Records held by a DNS-type distributed database system so that instead of the parsed telephone number being resolved to an URI which is then used to access the required resource, the parsed telephone number is directly resolved to the required service resource item. The mechanism employed in this process is exactly as already described for resolving a parsed telephone number into an URI. The DNS-type distributed database system used for this would

US 8,170,008 B2

31

preferably be one accessible over the Internet or the DNS itself so as to provide access to the service resource items for Internet users as well as for the service control subsystem of the PSTN (in the same manner as described above with reference to FIG. 18, the DNS-type servers holding the service resource items may be accessible to the service control subsystem by a network other than the Internet). Whilst the placing of service resource items in RRs held on DNS-type servers may not be suitable for all types of service resource items, it is suitable for items such as telephone numbers that do not change frequently. Thus, a suitable usage is to provide number portability; in this case, a dialed personal number triggers a lookup in the DNS-type system with all or part of the personal number being first parsed and then applied to the DNS type system to return a current number for call routing. All dialed numbers could be treated as personal numbers or simply a subset of such numbers, this subset comprising numbers that are readily identifiable as personal numbers by, for example, local lookup at an SSP or the presence of a predetermined leading digit string. The general concept of parsing a telephone number (or similar number) in whole or in part to form a domain name for resolution in a DNS-type distributed database system can be used for the retrieval of other items of information besides URIs and service resource items.

Feedback Mechanisms. In discussing the WWW-based freephone arrangement of FIG. 17, it was mentioned that user A could be supplied with feedback on the likely length of waiting time before A would be called back. This is one example of using the Internet to provide a feedback path for a potential or actual telephone user. Another example was provided in relation to FIG. 16 where the progress of call setup was reported back by the call setup gateway to user A's terminal. In fact, generally where a user is known to be using a terminal actively on the Internet the opportunity arises to provide the user with feedback on the progress of call setup through the telephone system. In order to do this, it is of course necessary to ensure that the feedback can be passed to the appropriate application running on terminal A and this will generally require the application to have made appropriate linking information available. As well as call setup progress information, other information can also be feedback for example during a call holding period. Thus, for example, a special server can be provided on the Internet holding multimedia clips or even videos that could be output to user A during a call holding period.

In the described arrangements, the servers 51 have held service resource items concerned primarily with call setup control. It may be noted that in a somewhat different application, Internet servers could be arranged to hold data that could be accessed from the telephone system in response to a user-initiated telephone request and returned to that telephone user. Such a service would be provided, for example, in response to an SSP triggering a service request upon a particular telephone number being input, the service request prompting an SCP to cause an intelligent peripheral to access a particular Internet server (not necessarily an HTTP server) and retrieve the required data for return to the calling party. The intelligent peripheral may include a text-to-voice converter for replaying the data vocally to the user.

One further feedback process is also worthy of note, in this case in relation to service resource items themselves. By way of example, a telephone user G may subscribe to a service by which calls passed through to G's telephone are to be separated by a minimum of X minutes, X being user settable. To implement this service, G has a phone page on a server 51 that includes a "busy" status indication. Upon termination of a

32

successful call to G, G's local SSP triggers the sending of a message by the associated SCP over the Internet to G's phone page. This message causes G's busy indication to be set to indicate that G is busy; the message also starts a timer which times out after a period X and causes the busy status indication to be reset. A call attempt to G will either be rejected at G's SSP because G's line is genuinely busy or will trigger the SSP to enquire via the SCP whether G's phone-page busy status indication is set. If the busy status indication is set (which it will be during the period X following termination of a successful call) the call attempt is rejected whereas if the busy status indication is in its reset condition, the call attempt is allowed to proceed.

By placing the busy status indication mechanism on G's phone page, it is possible to arrange for G to be able to easily change the value of X.

More General Variants. Whilst the service control subsystem of the PSTN has been embodied as an SCP in the foregoing examples, it will be appreciated that the functionality of the service control subsystem could be provided as part of an SSP or in an associated adjunct. Furthermore, the triggering of service requests can be effected by equipment other than SSPs, for example by intercept boxes inserted in the SS7 signalling links.

It will be appreciated that the term "Internet" is to be understood to include not only the current specification of the TCP/IP protocols used for the Internet and the current addressing scheme, but also evolutions of these features such as may be needed to deal with isochronous media. Furthermore, references to the WWW and the HTTP protocol should equally be understood to encompass their evolved descendants.

The present invention can also be applied to telephone systems other than just PSTNs, for example to PLMNs and other mobile networks, and to private systems using PABXs. In this latter case, a LAN or campus-wide computer network serving generally the same internal users as the PABX, will take the role of the Internet in the described embodiments.

Furthermore, the present invention has application where any switched telecommunication system (for example, a broadband ATM system) requires service control and a computer network can be used for the delivery of service resources to the service control subsystem of the telecommunication system.

The invention claimed is:

1. A method, comprising:

receiving, over a switched telecommunication system, a request;

determining, responsive to the request, a call destination using domain name system signaling; and

initiating a call through the switched telecommunication system between a calling party and the call destination that was determined as a result of said domain name system signaling.

2. The method of claim 1, wherein said initiating the call comprises using SS7 signaling to initiate the call between the calling party and the call destination.

3. The method of claim 1, wherein the switched telecommunication system comprises at least one of a public switched telephone network (PSTN) and a mobile network.

4. The method of claim 1, wherein said using comprises at least one of: (1) using domain name system signaling in the Internet, and (2) using domain name system signaling in a private system.

5. The method of claim 1, wherein the request indicates a first telephone number, and the call destination determined as

US 8,170,008 B2

33

a result of said domain name system signaling is associated with a second telephone number.

6. The method of claim 1, wherein the call destination is a destination in the switched telecommunication system.

7. The method of claim 1, wherein the call destination is a destination in a communication system coupled to the switched telecommunication system.

8. The method of claim 1, wherein said using domain name system signaling comprises:

5 sending a query to a database system; and
receiving from the database system a response to the query, the response indicating the call destination.

9. The method of claim 8, wherein the database system comprises the domain name system of the Internet.

10. The method of claim 8, wherein the database system is independent of the domain name system of the Internet.

11. The method of claim 1, further comprising using the request to determine a domain name, wherein said using domain name system signaling comprises:

20 sending a query to a database system, the query comprising the domain name; and
receiving from the database system a response to the query, the response indicating the call destination, wherein the response comprises a uniform resource name, and said determining the call destination comprises using the uniform resource name to determine the call destination.

12. The method of claim 1, wherein the call is through both the switched telecommunication system and another telecommunication system.

13. A method, comprising:

receiving, over a switched telecommunication system, an indication of a called party;
determining, responsive to the indication of the called party, a call destination associated with the called party using domain name system signaling; and
initiating a call through the switched telecommunication system between a calling party and the call destination that was determined as a result of said domain name system signaling.

14. The method of claim 13, wherein said initiating comprises using SS7 signaling.

15. The method of claim 14, wherein said initiating comprises also using at least one of TCAP signaling and ISUP signaling.

16. The method of claim 13, wherein using domain name system signaling comprises sending a query.

17. The method of claim 13, wherein using comprises at least one of: (1) using domain name system signaling in the Internet, and (2) using domain name system signaling in a private system.

18. The method of claim 13, further comprising using the indication of the called party to determine a domain name, wherein said using domain name system signaling comprises:

50 sending a query to a database system, the query comprising the domain name; and
receiving from the database system a response to the query, the response indicating the call destination, wherein the response comprises a uniform resource name, and said determining the call destination comprises using the uniform resource name to determine the call destination.

34

19. The method of claim 13, wherein the call is through both the switched telecommunication system and another telecommunication system.

20. A method, comprising:

receiving, over a switched telecommunication system, a request;

determining, responsive to the request, a record using domain name system signaling;

using the record to determine an indication of a call destination; and

initiating a call through the switched telecommunication system between a calling party and the call destination that was determined using said record that was determined using domain name system signaling.

21. The method of claim 20, wherein said initiating comprises using SS7 signaling to set up the call.

22. The method of claim 21, wherein said initiating comprises also using at least one of TCAP signaling and ISUP signaling.

23. The method of claim 20, wherein the call destination is a destination in the switched telecommunication system.

24. The method of claim 20, wherein the call destination is a destination in a communication system coupled to the switched telecommunication system.

25. The method of claim 20, wherein said determining the record comprises:

sending a query; and
receiving the record in response to the query.

26. The method of claim 20, further comprising using the request to determine a domain name, wherein said determining the record comprises:

30 sending a query to a database system, the query comprising the domain name; and
receiving the record in response to the query, wherein the record comprises a uniform resource name, and said using the record comprises using the uniform resource name to determine the indication of the call destination.

27. A method, comprising:

receiving a request from a first party;
determining by a computer, responsive to the request, an identifier of a second party using domain name system signaling;

45 setting up a call through the switched telecommunication system between the first party and the second party that was determined as a result of said domain name system signaling.

28. The method of claim 27, further comprising using the request to determine a domain name, wherein said using domain name system signaling comprises:

50 sending a query to a database system, the query comprising the domain name; and
receiving from the database system a response to the query, wherein the response comprises a uniform resource name, and said determining the indication of the second party comprises using the uniform resource name to determine the indication of the second party.

29. The method of claim 27, wherein the call is through both the switched telecommunication system and another telecommunication system.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 8,170,008 B2
APPLICATION NO. : 12/389843
DATED : May 1, 2012
INVENTOR(S) : Colin Low et al.

Page 1 of 2

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Title Page

Item 56 Under U.S. Patent Documents,

Please insert the following considered reference:

--US 20020167940 11/2002 Low, et al.--

Under Other Publications,

Please insert the following considered references:

--VoIP Watch, "All About Voice Over IP and Internet Based Telephony", April 19, 2007,
2 pages.

Kahane, et al., "Call Management Agent System Specification", VoIP-0017, August 15,
1996, 36 pages.

Sebestyen, "What is the Position of Q.2, Q.3/15 on Internet Telephony for the IMTC
VoIP Forum Meeting in Seattle", ITU Telecommunication Standardization Sector, Study
Group 15, Q.2&3/15 Rapporteur Meeting, December 17-20, 1996, 5 pages.

Petrack, "Backup and Alternate Gatekeepers in H.323 for High Availability, Reliability
and Scalability", ITU Telecommunication Standardization Sector, Study Group 16, Ex.
Q.2&3/15 Rapporteur Meeting, February 1-21, 1997, 7 pages.

Chapter 5: The CMA System, VoIP Forum, October 11, 1996, 38 pages.

CMA IPR Additions, November 17, 2008, 3 pages.

Rosenberg, et al., "Internet Telephony Gateway Location", In Proc. Of IEEE INFOCOM,
March 1998, 9 pages.

IMTC Voice Over IP Forum Service Interoperability Implementation Agreement, Draft
0.91, VOIP97-008, January 13, 1997, 70 pages.

IMTC VoIP Forum Contribution, Proposal for Use of LDAP for Internet Telephony
Directory Services, March 11, 1997, 8 pages.

Dynamic Address Mappings for Basic Call Management within the VoIP IA 1.0, June
16, 1997, 5 pages.

Signed and Sealed this
Nineteenth Day of February, 2013



Teresa Stanek Rea
Acting Director of the United States Patent and Trademark Office

CERTIFICATE OF CORRECTION (continued)

Page 2 of 2

U.S. Pat. No. 8,170,008 B2

“Prior Art on Verizon’s VoIP Patents,”
<http://www.techdirt.com/articles/20070417/005814.shtml>, April 17, 2007, 9 pages.
 “Another Day, Another Set of Prior Art Discovered Against Verizon’s VoIP Patents,”
<http://www.techdirt.com/articles/20070423/094554.shtml>, April 23, 2007, 5 pages.
 Verizon Gets Smacked Down for its VoIP Patent Suing Spree,
www.techdirt.com/articles/20081006/2246322473.shtml, October 6, 2008, 3 pages.
 Shockey, “ENUM: Phone Numbers Meet the Net”,
<http://www.callcentermagazine.com/article/CTM20010618S0010>, June 18, 2001, 4 pages.
 Infrastructure ENUM, http://www.circleid.com/posts/infrastructure_enum/, April 5, 2007, 7
 pages.
 TPC.INT, <http://osdir.com/ml/ietf.fax/1996-12/msg00161.html>, December 19, 1996,
 2 pages.
 Decker, “New Internet/Telecom Product/Technology Wanted”,
[http://groups.google.com/group/comp.dcom.telecom/browse_thread/thread/1227ba19866a342](http://groups.google.com/group/comp.dcom.telecom/browse_thread/thread/1227ba19866a342/2...)
 /2...>, September 22, 1995, 3 pages.
 Michalski, “The Internet as Communications Lab”, Release 1.0, 2-94, February 18, 1994, 28
 pages.
 Free World Dialup and Verizon’s patent on “name translation”, The Jeff Pulver Blog,
<http://pulverblog.pulver.com/archives/006846.html>, April 23, 2007, 4 pages.
 Rose, et al., “An Experiment in Remote Printing”, RFC1486,
<http://tools.ietf.org/html/rfc1486>, July 1993, 14 pages.
 Alvestrand, et al., “Mapping between X.400 and RFC-822 Message Bodies”, RFC1495,
<http://tools.ietf.org/html/rfc1495>, August 1993, 12 pages.
 Rose, “Principles of Operation for the TPC.INT Subdomain: Radio Paging -- Technical
 Procedures”, RFC1569, <http://tool.ietf.org/html/rfc1569>, January 1994, 7 pages.
 Klensin, et al., “SMTP Service Extensions”, RFC1651, <http://tools.ietf.org/html/rfc1651>,
 July 1994, 12 pages.
 Klensin, et al., “SMTP Service Extension for 8bit-MIMEtransport”, RFC1652,
<http://tool.ietf.org/html/rfc1652>, July 1994, 7 pages.
 Rose, “Principles of Operation for the TPC.INT Subdomain: Radio paging -- Technical
 Procedures”, RFC1703, <http://tools.ietf.org/html/rfc1703>, October 1994, 10 pages.
 Telecom Digest, September 21, 1995, Volume 15, Issue 400,
<http://massis.lcs.mit.edu/archives/back.issues/1995.volume.15/vol15.iss351-400>, 5 pages.
 Validity of Verizon’s VoIP Patents Challenged – PC World,
http://www.pcworld.com/article/130864/validity_of_verizons_voip_patents_challenged.html
 April 18, 2007, 4 pages.--

(12) **United States Patent**
Low et al.

(10) **Patent No.:** **US 8,204,046 B2**
(45) **Date of Patent:** ***Jun. 19, 2012**

(54) **METHOD AND APPARATUS FOR ACCESSING SERVICE RESOURCE ITEMS THAT ARE FOR USE IN A TELECOMMUNICATIONS SYSTEM**

(75) Inventors: **Colin Low**, Gloucestershire (GB);
Andrew Franklin Seaborne, Bristol (GB); **Nicolas Bouthors**, Les Bealires (FR)

(73) Assignee: **Comcast IP Holdings I, LLC**,
Wilmington, DE (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 72 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **12/852,787**

(22) Filed: **Aug. 9, 2010**

(65) **Prior Publication Data**

US 2010/0303065 A1 Dec. 2, 2010

Related U.S. Application Data

(60) Continuation of application No. 11/066,880, filed on Feb. 25, 2005, now Pat. No. 7,804,816, which is a division of application No. 10/052,285, filed on Jan. 18, 2002, now Pat. No. 7,012,916, which is a division of application No. 09/077,795, filed as application No. PCT/GB96/03055 on Dec. 11, 1996, now Pat. No. 6,466,570.

(30) **Foreign Application Priority Data**

Dec. 11, 1995 (GB) 9525190.6
Dec. 22, 1995 (EP) 95410148
Feb. 20, 1996 (GB) 9603582.9

(51) **Int. Cl.**

H04L 12/66 (2006.01)
G06F 15/16 (2006.01)

G06F 7/00 (2006.01)
G06F 17/30 (2006.01)
(52) **U.S. Cl.** **370/352**; 709/206; 709/219; 707/755; 707/769

(58) **Field of Classification Search** None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,054,756 A 10/1977 Comella et al.
(Continued)

FOREIGN PATENT DOCUMENTS

EP 0559979 9/1993
(Continued)

OTHER PUBLICATIONS

IMTC VoIP Forum Contribution, Proposal for Use of LDAP for Internet Telephony Directory Services, Mar. 11, 1997, 8 pages.

(Continued)

Primary Examiner — Andrew Chriss


(74) *Attorney, Agent, or Firm* — Banner & Witcoff, Ltd.

(57)

ABSTRACT

Service resource items for use in call setup in a telephone system are held on servers that are connected to a computer network which is logically distinct from the telephone system infrastructure; this computer network may, for example, make use of the Internet. Each service item is locatable on the network at a corresponding URI and is associated with a particular telephone number. A mapping is provided between telephone numbers and the URIs of associated service resource items. When it is desired to access a service resource item associated with a particular telephone number, this mapping is used to retrieve the corresponding URI which is then used to access the desired service resource item.

115 Claims, 14 Drawing Sheets



US 8,204,046 B2

Page 2

U.S. PATENT DOCUMENTS					
4,100,377	A	7/1978	Flanagan	5,440,620	A 8/1995 Slusky
4,201,891	A	5/1980	Lawrence et al.	5,448,633	A 9/1995 Jamaledin et al.
4,313,035	A	1/1982	Jordan et al.	5,450,411	A 9/1995 Heil
4,611,094	A	9/1986	Asmuth et al.	5,452,289	A 9/1995 Sharma et al.
4,611,096	A	9/1986	Asmuth et al.	5,452,350	A 9/1995 Reynolds et al.
4,653,045	A	3/1987	Stanley et al.	5,455,821	A 10/1995 Schaeffer et al.
4,659,877	A	4/1987	Dorsey et al.	5,459,775	A 10/1995 Isono et al.
4,679,190	A	7/1987	Dias et al.	5,461,611	A 10/1995 Drake, Jr. et al.
4,685,125	A	8/1987	Zave	5,463,677	A 10/1995 Bash et al.
4,713,806	A	12/1987	Oberlander et al.	5,469,496	A 11/1995 Emery et al.
4,718,005	A	1/1988	Feigenbaum et al.	5,469,497	A 11/1995 Pierce et al.
4,734,931	A	3/1988	Bourg et al.	5,473,608	A 12/1995 Gagne et al.
4,747,130	A	5/1988	Ho	5,475,748	A 12/1995 Jones
4,771,425	A	9/1988	Baran et al.	5,475,817	A 12/1995 Waldo et al.
4,788,718	A	11/1988	McNabb et al.	5,477,531	A 12/1995 McKee et al.
4,821,034	A	4/1989	Anderson et al.	5,479,494	A 12/1995 Clitherow
4,872,157	A	10/1989	Hemmady et al.	5,481,603	A 1/1996 Gutierrez et al.
4,893,302	A	1/1990	Hemmady et al.	5,483,586	A 1/1996 Sussman
4,897,874	A	1/1990	Lidinsky et al.	5,483,587	A 1/1996 Hogan et al.
4,899,373	A	2/1990	Lee et al.	5,487,111	A 1/1996 Slusky
4,933,931	A	6/1990	Kokubo	5,495,521	A 2/1996 Rangachar
4,969,184	A	11/1990	Gordon et al.	5,500,889	A 3/1996 Baker et al.
4,979,206	A	12/1990	Padden et al.	5,500,929	A * 3/1996 Dickinson 715/853
4,996,707	A	2/1991	O'Malley et al.	5,506,887	A 4/1996 Emery et al.
5,025,254	A	6/1991	Hess	5,506,893	A 4/1996 Buscher et al.
5,029,196	A	7/1991	Morganstein	5,511,111	A 4/1996 Serbetcioglu et al.
5,065,393	A	11/1991	Sibbitt et al.	5,513,127	A 4/1996 Gard et al.
5,068,888	A	11/1991	Scherk et al.	5,521,719	A 5/1996 Yamada
5,107,492	A	4/1992	Roux et al.	5,521,924	A 5/1996 Kakuma et al.
5,113,499	A	5/1992	Ankney et al.	5,524,137	A 6/1996 Rhee
5,115,495	A	5/1992	Tsuchiya et al.	5,526,353	A 6/1996 Henley et al.
5,146,488	A	9/1992	Okada et al.	5,526,416	A 6/1996 Dezonno et al.
5,146,491	A	9/1992	Silver et al.	5,526,489	A 6/1996 Nilakantan et al.
5,157,390	A	10/1992	Yoshie et al.	5,530,744	A 6/1996 Charalambous et al.
5,157,662	A	10/1992	Tadamura et al.	5,530,852	A 6/1996 Meske, Jr. et al.
5,159,592	A	10/1992	Perkins	5,533,115	A 7/1996 Hollenbach et al.
5,159,624	A	10/1992	Makita	5,539,884	A 7/1996 Robrock, II
5,164,938	A	11/1992	Jurkevich et al.	5,539,886	A 7/1996 Aldred et al.
5,185,860	A	2/1993	Wu	5,541,917	A 7/1996 Farris
5,193,110	A	3/1993	Jones et al.	5,541,927	A 7/1996 Kristol et al.
5,195,085	A	3/1993	Bertsch et al.	5,541,930	A 7/1996 Klingman
5,195,086	A	3/1993	Baumgartner et al.	5,546,452	A 8/1996 Andrews et al.
5,204,894	A	4/1993	Darden	5,551,025	A 8/1996 O'Reilly et al.
5,206,901	A	4/1993	Harlow et al.	5,551,035	A 8/1996 Arnold et al.
5,223,699	A	6/1993	Flynn et al.	5,559,721	A 9/1996 Ishii
5,226,075	A	7/1993	Funk et al.	5,559,871	A 9/1996 Smith
5,241,588	A	8/1993	Babson, III et al.	5,561,670	A 10/1996 Hoffer et al.
5,247,571	A	9/1993	Kay et al.	5,563,882	A 10/1996 Bruno et al.
5,260,986	A	11/1993	Pershan	5,570,410	A 10/1996 Hooshiari
5,272,749	A	12/1993	Masek	5,575,961	A 11/1996 Kuwabara et al.
5,287,199	A	2/1994	Zoccolillo	5,579,472	A 11/1996 Keyworth, II et al.
5,311,583	A	5/1994	Friedes et al.	5,583,920	A 12/1996 Wheeler, Jr.
5,317,566	A	5/1994	Joshi	5,586,177	A 12/1996 Farris et al.
5,325,421	A	6/1994	Hou et al.	5,590,127	A 12/1996 Bales et al.
5,327,486	A	7/1994	Wolff et al.	5,590,181	A 12/1996 Hogan et al.
5,329,520	A	7/1994	Richardson	5,594,769	A 1/1997 Pellegrino et al.
5,329,578	A	7/1994	Brennan et al.	5,594,789	A 1/1997 Seazholtz et al.
5,333,266	A *	7/1994	Boaz et al. 709/206	5,603,084	A 2/1997 Henry, Jr. et al.
5,347,633	A	9/1994	Ashfield et al.	5,603,085	A 2/1997 Shedlo
5,351,276	A	9/1994	Doll, Jr. et al.	5,604,682	A 2/1997 McLaughlin et al.
5,351,286	A	9/1994	Nici	5,604,737	A 2/1997 Iwami et al.
5,353,331	A	10/1994	Emery et al.	5,608,446	A 3/1997 Carr et al.
5,361,256	A	11/1994	Doeringer et al.	5,608,786	A 3/1997 Gordon
5,390,335	A	2/1995	Stephan et al.	5,610,910	A 3/1997 Focsaneanu et al.
5,392,344	A	2/1995	Ash et al.	5,610,972	A 3/1997 Emery et al.
5,396,542	A	3/1995	Alger et al.	5,610,976	A 3/1997 Uota et al.
5,402,478	A	3/1995	Hluchyj et al.	5,617,422	A 4/1997 Litzenberger et al.
5,406,557	A	4/1995	Baudoin	5,617,540	A 4/1997 Civanlar et al.
5,418,844	A	5/1995	Morrissey et al.	5,619,555	A 4/1997 Fenton et al.
5,422,940	A	6/1995	Endo et al.	5,619,557	A 4/1997 Van Berkum
5,423,003	A	6/1995	Berteau	5,619,562	A 4/1997 Maurer et al.
5,425,028	A	6/1995	Britton et al.	5,623,601	A 4/1997 Vu
5,425,085	A	6/1995	Weinberger et al.	5,625,404	A 4/1997 Grady et al.
5,425,091	A	6/1995	Josephs	5,625,407	A 4/1997 Biggs et al.
5,428,608	A	6/1995	Freeman et al.	5,625,675	A 4/1997 Katsumaru et al.
5,436,963	A	7/1995	Fitzpatrick et al.	5,625,677	A 4/1997 Feiertag et al.
5,438,568	A	8/1995	Weisser, Jr.	5,625,682	A 4/1997 Gray et al.
5,440,563	A	8/1995	Isidoro et al.	5,627,886	A 5/1997 Bowman
				5,629,978	A 5/1997 Blumhardt et al.

US 8,204,046 B2

Page 3

5,633,916 A	5/1997	Goldhagen et al.	5,812,795 A	9/1998	Horovitz et al.
5,634,012 A	5/1997	Stefik et al.	5,812,865 A	9/1998	Theimer et al.
5,636,216 A	6/1997	Fox et al.	5,818,921 A	10/1998	Vander Meiden et al.
5,646,982 A	7/1997	Hogan et al.	5,825,772 A	10/1998	Dobbins et al.
5,651,006 A	7/1997	Fujino et al.	5,825,869 A	10/1998	Brooks et al.
5,652,787 A	7/1997	O'Kelly	5,826,268 A	10/1998	Schaefer et al.
5,654,957 A	8/1997	Koyama	5,828,737 A	10/1998	Sawyer
5,657,250 A	8/1997	Park et al.	5,828,844 A	10/1998	Civanlar et al.
5,661,197 A	8/1997	Villiger et al.	5,832,197 A	11/1998	Houji
5,661,782 A	8/1997	Bartholomew et al.	5,835,723 A	11/1998	Andrews et al.
5,661,790 A	8/1997	Hsu	5,838,682 A	11/1998	Dekelbaum et al.
5,664,005 A	9/1997	Emery et al.	5,838,686 A	11/1998	Ozkan
5,664,102 A	9/1997	Faynberg	5,838,768 A	11/1998	Sumar et al.
5,668,857 A	9/1997	McHale	5,838,970 A	11/1998	Thomas
5,673,263 A	9/1997	Basso et al.	5,845,267 A	12/1998	Ronen
5,675,507 A	10/1997	Bobo, II	5,848,143 A	12/1998	Andrews et al.
5,675,741 A	10/1997	Aggarwal et al.	5,850,433 A	12/1998	Rondeau
5,680,116 A	10/1997	Hashimoto et al.	5,864,604 A	1/1999	Moen et al.
5,680,392 A	10/1997	Semaan	5,864,610 A	1/1999	Ronen
5,687,167 A	11/1997	Bertin et al.	5,864,670 A	1/1999	Hayashi et al.
5,689,550 A	11/1997	Garson et al.	5,867,495 A	2/1999	Elliott et al.
5,689,553 A	11/1997	Ahuja et al.	5,867,566 A	2/1999	Hogan et al.
5,692,039 A	11/1997	Brankley et al.	5,867,571 A	2/1999	Borchering
5,694,463 A	12/1997	Christie et al.	5,870,454 A	2/1999	Dahlen
5,699,089 A	12/1997	Murray	5,870,561 A	2/1999	Jarvis et al.
5,699,352 A	12/1997	Kriete et al.	5,873,077 A	2/1999	Kanoh et al.
5,701,295 A	12/1997	Bales et al.	5,878,126 A	3/1999	Velamuri et al.
5,701,465 A	12/1997	Baugher et al.	5,878,128 A	3/1999	Kantola
5,703,935 A	12/1997	Raissy et al.	5,878,212 A	3/1999	Civanlar et al.
5,703,940 A	12/1997	Sattar et al.	5,881,134 A	3/1999	Foster et al.
5,703,942 A	12/1997	Pinard et al.	5,883,891 A	3/1999	Williams et al.
5,708,780 A *	1/1998	Levergood et al. 709/229	5,883,894 A	3/1999	Patel et al.
5,710,884 A	1/1998	Dedrick	5,884,032 A	3/1999	Bateman et al.
5,712,903 A	1/1998	Bartholomew et al.	5,889,774 A	3/1999	Mirashrafi et al.
5,712,906 A	1/1998	Grady et al.	5,889,777 A	3/1999	Miyao et al.
5,712,907 A	1/1998	Wegner et al.	5,892,754 A	4/1999	Kompella et al.
5,712,908 A	1/1998	Brinkman et al.	5,903,558 A	5/1999	Jones et al.
5,719,854 A	2/1998	Choudhury et al.	5,903,559 A	5/1999	Acharya et al.
5,724,355 A	3/1998	Bruno et al.	5,903,735 A	5/1999	Kidder et al.
5,724,406 A	3/1998	Juster	5,907,547 A	5/1999	Foladare et al.
5,724,412 A	3/1998	Srinivasan	5,909,430 A	6/1999	Reaves
5,726,984 A	3/1998	Kubler et al.	5,915,008 A	6/1999	Dulman
5,727,129 A	3/1998	Barrett et al.	5,915,012 A	6/1999	Miloslavsky
5,729,544 A	3/1998	Lev et al.	5,917,817 A	6/1999	Dunn et al.
5,729,599 A	3/1998	Plomondon et al.	5,923,659 A	7/1999	Curry et al.
5,732,078 A	3/1998	Arango	5,930,348 A	7/1999	Regnier et al.
5,732,219 A	3/1998	Blumer et al.	5,933,412 A	8/1999	Choudhury et al.
5,737,333 A	4/1998	Civanlar et al.	5,933,425 A	8/1999	Iwata
5,737,395 A	4/1998	Irribarren	5,933,486 A	8/1999	Norby et al.
5,737,414 A	4/1998	Walker et al.	5,936,958 A	8/1999	Soumiya et al.
5,742,668 A	4/1998	Pepe et al.	5,940,479 A	8/1999	Guy et al.
5,742,905 A	4/1998	Pepe et al.	5,940,616 A	8/1999	Wang
5,744,533 A	4/1998	Iwamoto et al.	5,944,795 A	8/1999	Civanlar
5,745,694 A	4/1998	Egawa et al.	5,949,869 A	9/1999	Sink et al.
5,748,468 A	5/1998	Notenboom et al.	5,949,977 A	9/1999	Hernandez
5,751,961 A	5/1998	Smyk	5,953,322 A	9/1999	Kimball
5,754,639 A	5/1998	Flockhart et al.	5,953,338 A	9/1999	Ma et al.
5,757,784 A	5/1998	Liebowitz et al.	5,953,392 A	9/1999	Rhie et al.
5,758,281 A	5/1998	Emery et al.	5,953,504 A	9/1999	Sokal et al.
5,761,290 A	6/1998	Farris et al.	5,956,391 A	9/1999	Melen et al.
5,774,530 A	6/1998	Montgomery et al.	5,956,482 A	9/1999	Agraharam et al.
5,774,533 A	6/1998	Patel	5,956,697 A	9/1999	Usui
5,774,670 A *	6/1998	Montulli 709/227	5,958,016 A	9/1999	Chang et al.
5,781,550 A	7/1998	Templin et al.	5,960,416 A	9/1999	Block
5,781,624 A	7/1998	Mitra et al.	5,966,427 A	10/1999	Shaffer et al.
5,784,617 A	7/1998	Greenstein et al.	5,966,434 A	10/1999	Schafer et al.
5,787,160 A	7/1998	Chaney et al.	5,970,065 A	10/1999	Miloslavsky
5,790,172 A	8/1998	Imanaka	5,978,773 A *	11/1999	Hudetz et al. 705/23
5,790,806 A	8/1998	Koperda	5,978,840 A	11/1999	Nguyen et al.
5,793,762 A	8/1998	Penners et al.	5,987,108 A	11/1999	Jagadish et al.
5,796,790 A	8/1998	Brunner	5,995,503 A	11/1999	Crawley et al.
5,799,063 A	8/1998	Krane	5,995,606 A	11/1999	Civanlar et al.
5,799,317 A	8/1998	He et al.	5,999,524 A	12/1999	Corbalis et al.
5,802,146 A	9/1998	Dulman	6,011,975 A	1/2000	Emery et al.
5,802,502 A	9/1998	Gell et al.	6,014,379 A	1/2000	White et al.
5,802,510 A	9/1998	Jones	6,014,437 A	1/2000	Acker et al.
5,809,129 A	9/1998	Andersson et al.	6,014,687 A	1/2000	Watanabe et al.
5,812,656 A	9/1998	Garland et al.	6,016,307 A	1/2000	Kaplan et al.
5,812,776 A	9/1998	Gifford et al.	6,018,567 A	1/2000	Dulman

US 8,204,046 B2

Page 4

6,021,126	A	2/2000	White et al.	6,466,570	B1	10/2002	Low et al.
6,021,263	A	2/2000	Kujoory et al.	6,466,984	B1	10/2002	Naveh et al.
6,023,724	A	2/2000	Bhatia et al.	6,470,081	B1	10/2002	Sbisa et al.
6,026,441	A	2/2000	Ronen	6,498,788	B1	12/2002	Emilsson et al.
6,029,203	A	2/2000	Bhatia et al.	6,519,254	B1	2/2003	Chuah et al.
6,035,028	A	3/2000	Ward et al.	6,529,516	B1	3/2003	Parzych
6,041,109	A	3/2000	Cardy et al.	6,556,659	B1	4/2003	Bowman-Amuah
6,041,117	A	3/2000	Androski et al.	6,584,093	B1	6/2003	Salama et al.
6,044,081	A	3/2000	Bell et al.	6,587,433	B1	7/2003	Borella et al.
6,047,051	A	4/2000	Ginzboorg et al.	6,594,254	B1	7/2003	Kelly
6,047,326	A	4/2000	Kilkki	6,600,733	B2	7/2003	Deng
6,049,545	A	4/2000	Stephenson et al.	6,636,485	B1	10/2003	Fijolek et al.
6,052,445	A	4/2000	Bashoura et al.	6,654,357	B1	11/2003	Wiedeman
6,052,450	A	4/2000	Allison et al.	6,671,285	B1	12/2003	Kirkby et al.
6,058,113	A	5/2000	Chang	6,678,264	B1	1/2004	Gibson
6,067,516	A	5/2000	Levay et al.	6,690,788	B1	2/2004	Bauer et al.
6,069,890	A	5/2000	White et al.	6,718,380	B1	4/2004	Mohaban et al.
6,073,160	A	6/2000	Grantham et al.	6,754,181	B1	6/2004	Elliott et al.
6,078,582	A	6/2000	Curry et al.	6,768,793	B1	7/2004	Sbisa et al.
6,084,956	A	7/2000	Turner et al.	6,775,519	B1	8/2004	Wiedeman et al.
6,088,358	A	7/2000	Tomita et al.	6,810,033	B2	10/2004	Derks
6,088,431	A	7/2000	LaDue	6,882,643	B1	4/2005	Mauger et al.
6,097,722	A	8/2000	Graham et al.	6,891,945	B2	5/2005	Porter et al.
6,097,804	A	8/2000	Gilbert et al.	6,983,040	B1	1/2006	Kalmanek, Jr. et al.
6,098,094	A	8/2000	Barnhouse et al.	7,013,001	B1	3/2006	Felger
6,101,182	A	8/2000	Sistanizadeh et al.	7,024,450	B1	4/2006	Deo et al.
6,104,704	A	8/2000	Buhler et al.	7,032,022	B1	4/2006	Shanumgam et al.
6,108,314	A	8/2000	Jones et al.	7,061,923	B2	6/2006	Dugan et al.
6,108,704	A	8/2000	Hutton et al.	7,065,571	B2	6/2006	Schweitzer et al.
6,115,737	A	9/2000	Ely et al.	7,088,705	B2	8/2006	Curry et al.
6,125,113	A	9/2000	Farris et al.	7,170,887	B2	1/2007	Rosenberg
6,131,095	A	10/2000	Low et al.	7,185,073	B1	2/2007	Gai et al.
6,134,433	A	10/2000	Joong et al.	7,236,575	B2	6/2007	Kim et al.
6,137,777	A	10/2000	Vaid et al.	7,286,562	B1	10/2007	Vargo et al.
6,141,413	A	10/2000	Waldner et al.	2001/0039576	A1	11/2001	Kanada
6,141,686	A	10/2000	Jackowski et al.	2001/0040885	A1	11/2001	Jonas et al.
6,144,647	A	11/2000	Lopez-Torres	2002/0041590	A1	4/2002	Donovan
6,144,661	A	11/2000	Katsube et al.	2002/0064149	A1	5/2002	Elliott et al.
6,144,667	A	11/2000	Doshi et al.	2002/0083166	A1	6/2002	Dugan et al.
6,151,319	A	11/2000	Dommety et al.	2002/0114324	A1	8/2002	Low et al.
6,154,777	A	11/2000	Ebrahim	2002/0167940	A1	11/2002	Low et al.
6,157,648	A	12/2000	Voit et al.	2003/0039210	A1	2/2003	Jin et al.
6,167,445	A	12/2000	Gai et al.	2003/0193933	A1	10/2003	Jonas et al.
6,169,735	B1	1/2001	Allen, Jr. et al.	2003/0198335	A1	10/2003	Porter et al.
6,169,748	B1	1/2001	Barbas et al.	2004/0005046	A1	1/2004	Deo et al.
6,175,618	B1	1/2001	Shah et al.	2004/0215787	A1	10/2004	Gibson et al.
6,181,690	B1	1/2001	Civanlar	2005/0021713	A1	1/2005	Dugan et al.
6,181,703	B1	1/2001	Christie et al.	2005/0191989	A1	9/2005	Plush et al.
6,185,198	B1	2/2001	LaDue	2006/0239254	A1	10/2006	Short et al.
6,188,677	B1	2/2001	Oyama et al.	2007/0050492	A1	3/2007	Jorgenson
6,195,355	B1	2/2001	Demizu	2007/0204036	A1	8/2007	Mohaban et al.
6,198,738	B1	3/2001	Chang et al.	2007/0263644	A1	11/2007	Christie et al.
6,205,148	B1	3/2001	Takahashi et al.				
6,212,162	B1	4/2001	Horlin				
6,233,604	B1	5/2001	Van Horne et al.				
6,240,091	B1	5/2001	Ginzboorg et al.	EP	0654930	5/1995	
6,243,373	B1	6/2001	Turock	EP	0729281	A2	8/1996
6,252,869	B1	6/2001	Silverman	EP	0740445		10/1996
6,260,067	B1	7/2001	Barnhouse et al.	EP	0767568		4/1997
6,278,693	B1	8/2001	Aldred et al.	EP	0781016		6/1997
6,278,704	B1	8/2001	Creamer et al.	EP	0802690		10/1997
6,282,281	B1	8/2001	Low	EP	0812089		12/1997
6,298,120	B1	10/2001	Civanlar et al.	EP	0823809		2/1998
6,298,383	B1	10/2001	Gutman et al.	EP	0729281	A3	1/2000
6,308,216	B1	10/2001	Goldschmidt et al.	EP	1091528	A2	11/2001
6,310,873	B1	10/2001	Rainis et al.	EP	1091528	A3	11/2001
6,324,264	B1	11/2001	Wiener et al.	JP	09168051		6/1997
6,324,279	B1	11/2001	Kalmanek, Jr. et al.	JP	09168063		6/1997
6,347,085	B2	2/2002	Kelly	JP	09168064		6/1997
6,359,880	B1	3/2002	Curry et al.	JP	09168065		6/1997
6,381,321	B1	4/2002	Brown et al.	JP	09172459		6/1997
6,385,193	B1	5/2002	Civanlar et al.	JP	09172462		6/1997
6,418,461	B1	7/2002	Barnhouse et al.	WO	9107839		5/1991
6,418,467	B1	7/2002	Schweitzer et al.	WO	9325035		12/1993
6,434,624	B1	8/2002	Gai et al.	WO	9411813		5/1994
6,446,200	B1	9/2002	Ball et al.	WO	9423523		10/1994
6,449,650	B1	9/2002	Westfall et al.	WO	9522221		8/1995
6,456,617	B1	9/2002	Oda et al.	WO	9529564		11/1995
6,463,470	B1	10/2002	Mohaban et al.	WO	9609714		3/1996
				WO	9620448		7/1996

FOREIGN PATENT DOCUMENTS

US 8,204,046 B2

Page 5

WO	9620553	7/1996
WO	9632800	10/1996
WO	9634341	10/1996
WO	9638018	11/1996
WO	9714238	4/1997
WO	9716007	5/1997
WO	9720424	6/1997
WO	9722211	6/1997
WO	9723078	6/1997
WO	9726749	7/1997
WO	9728628	8/1997
WO	9732427	9/1997
WO	9733412	9/1997
WO	9823080	5/1998
WO	9834391 A2	8/1998
WO	9834391 A3	8/1998

OTHER PUBLICATIONS

Dynamic Address Mappings for Basic Call Management within the VoIP IA 1.0, Jun. 16, 1997, 5 pages.

"Prior Art on Verizon's VoIP Patents," <<http://www.techdirt.com/articles/20070417/005814.shtml>>, Apr. 17, 2007, 9 pages.

Another Day, Another Set of Prior Art Discovered Against Verizon's VoIP Patents, <<http://www.techdirt.com/articles/20070423/094554.shtml>>, Apr. 23, 2007, 5 pages.

Verizon Gets Smacked Down for Its VoIP Patent Suing Spree, <<http://www.techdirt.com/articles/20081006/2246322473.shtml>>, Oct. 6, 2008, 3 pages.

Shockey, "ENUM: Phone Numbers Meet the Net", <<http://www.callcentermagazine.com/article/CTM20010618S0010>>, Jun. 18, 2001, 4 pages.

Infrastructure ENUM, <http://www.circleid.com/posts/infrastructure_enum/>, Apr. 5, 2007, 7 pages.

TPC.INT: msg#00161, <<http://osdir.com/ml/ietf.fax/1996-12/msg00161.html>>, Dec. 19, 1996, 2 pages.

Decker, "New Internet/Telecom Product/Technology Wanted", <http://groups.google.com/group/comp.dcom.telecom/browse_thread/thread/1227ba19866a3426/2...>, Sep. 22, 1995, 3 pages.

Michalski, "The Internet as Communications Lab", Release 1.0, 2-94, Feb. 18, 1994, 28 pages.

Free World Dialup and Verizon's patent on "name translation", The Jeff Pulver Blog, <<http://pulverblog.pulver.com/archives/006846.html>>, Apr. 23, 2007, 4 pages.

Rose, et al., "An Experiment in Remote Printing", RFC1486, <<http://tools.ietf.org/html/rfc1486>>, Jul. 1993, 14 pages.

Alvestrand, et al., "Mapping between X.400 and RFC-822 Message Bodies", RFC1495, <<http://tools.ietf.org/html/rfc1495>>, Aug. 1993, 12 pages.

Rose, "Principles of Operation for the TPC.INT Subdomain: Radio Paging—Technical Procedures", RFC1569, <<http://tool.ietf.org/html/rfc1569>>, Jan. 1994, 7 pages.

Klensin, et al., "SMTP Service Extensions", RFC1651, <<http://tools.ietf.org/html/rfc1651>>, Jul. 1994, 12 pages.

Klensin, et al., "SMTP Service Extension for 8bit-MIMEtransport", RFC1652, <<http://tool.ietf.org/html/rfc1652>>, Jul. 1994, 7 pages.

Rose, "Principles of Operation for the TPC.INT Subdomain: Radio Paging—Technical Procedures", RFC1703, <<http://tools.ietf.org/html/rfc1703>>, Oct. 1994, 10 pages.

Telecom Digest, Sep. 21, 1995, vol. 15, Issue 400, <<http://massis.lcs.mit.edu/archives/back.issues/1995.volume.15/vol15.iss351-400>>, Sep. 21, 1995, 5 pages.

Validity of Verizon's VoIP Patents Challenged—PC World, <http://www.pcworld.com/article/130864/validity_of_verizons_voip_patents_challenged.html>, Apr. 18, 2007, 4 pages.

SG15 Plenary, May 28, 1996, regarding H.323, 82 pages.

RFC 1034, Nov. 1987, 41 pages.

RFC 1035, Nov. 1987, 49 pages.

SS7 Signaling Protocol, Message Transfer Part (MTP), Mar. 1996, 427 pages.

SS7 Signaling Protocol, Signalling Connection Control Part (SCCP), Oct. 1992, 220 pages.

Martersteck, et al., "The 5ESS Switching System: Introduction," AT&T Technical Journal, Jul. 1985, 10 pages.

Kahane, "Introduction to VocalTec's CMA System, Intelligent Switching for a New Era of Telecom", Nov. 7, 1996, pp. 1-18.

Karn, "Geek of the Week, Mobile IP Networking", Internet Talk Radio, Nov. 1993, 17 pages.

Yamada, et al., "A Study on IN Basic Call State Model for Packet Switched Network", IEEE, Apr. 1996, 18 pages.

ANSI, Signalling System No. 7 (SS7), General Information, ANSI, Jun. 1992, 78 pages.

Cohen, Specifications for the Network Voice Protocol (NVP), Jan. 29, 1976, 33 pages.

Everhart, et al., "New DNS RR Definitions", RFC 1183, Oct. 1990, 11 pages.

Schulzrinne, "Simple Conference Invitation Protocol", Internet Draft, Feb. 22, 1996, 19 pages.

Vocaltec's Telephony Gateway Software Captures 1996 Product of the Year Honors from Computer Telephony Magazine, Dec. 18, 1996, 2 pages.

Verjinski, "PHASE, A Portable Host Access System Environment", IEEE, Oct. 1989, 4 pages.

Schulzrinne, "Personal Mobility for Multimedia Services in the Internet", IDMS 1996, Mar. 1996, 18 pages.

Garrahan, et al., "Intelligent Network Overview", IEEE Communications Magazine, Mar. 1993, 7 pages.

Ash, et al., "Design & Optimization of Networks with Dynamic Routing", The Bell System Technical Journal, vol. 60, No. 8, Oct. 1981, pp. 1787-1820.

Graff, et al., "The 5ESS Switching System: Physical Design/Hardware," AT&T Technical Journal, Jul. 1985, 46 pages.

Bellcore Document GR-1428-CORE, CCS Network Interface Specification (CCSNIS), Supporting Toll-Free Service, A Module of CCSNIS, FR-905, May 1995, 93 pages.

Email from Dr. Schulzrinne regarding publication of SCIP, Oct. 10, 1995, 1 page.

GR-1280-CORE, Advanced Intelligent Network (AIN) Service Control Point (SCP) Generic Requirements, Telcordia Technologies Generic Requirements, Issue 1, Aug. 1993, 637 pages.

GR-2863-CORE, CCS Network Interface Specification (CCSNIS) Supporting Advanced Intelligent Network (AIN), Bellcore, Generic Requirements, Issue 2, Dec. 1995, 148 pages.

Office Action, U.S. Appl. No. 10/052,304 (U.S. Patent No. 7,206,305), mailed Jun. 19, 2006.

Office Action, U.S. Appl. No. 11/066,880 (pending), mailed Apr. 14, 2009.

Office Action, U.S. Appl. No. 11/066,880 (pending), mailed Nov. 23, 2009.

Gibson, "Web Foes Belong in 'Animal House'," PC Week, Aug. 19, 1996, 1 page.

VocalTech Homepage—2 pages, no earlier than Dec. 1996.

Supplementary European Search Report for EP00976942 dated Jul. 7, 2003.

International Search Report for PCT/US00/30448 dated Jan. 16, 2001, 3 pages.

Haugk, et al., "The 5ESS Switching System: Maintenance Capabilities," AT&T Technical Journal, Jul. 1985, 32 pages.

Basinger, et al., "The 5ESS Switching System: System Development Environment," AT&T Technical Journal, Jul. 1985, 18 pages.

Bauer, et al., "The 5ESS Switching System: System Test, First-Office Application, and Early Field Experience," AT&T Technical Journal, Jul. 1985, 20 pages.

Borum, et al., "The 5ESS Switching System: Hardware Design," AT&T Technical Journal, Jul. 1985, 21 pages.

Byrne, et al., "The 5ESS Switching System: Applications Planning," AT&T Technical Journal, Jul. 1985, 23 pages.

Carney, et al., "The 5ESS Switching System: Architectural Overview," AT&T Technical Journal, Jul. 1985, 18 pages.

Delatore, et al., "The 5ESS Switching System: Factory System Testing," AT&T Technical Journal, Jul. 1985, 22 pages.

Delatore, et al., "The 5ESS Switching System: Operational Software," AT&T Technical Journal, Jul. 1985, 28 pages.

Fuhrer, et al., "The 5ESS Switching System: Operations, Administration, and Maintenance Capabilities," AT&T Technical Journal, Jul. 1985, 14 pages.

US 8,204,046 B2

Page 6

- Office Action, U.S. Appl. No. 11/590,489 (issued), mailed Mar. 18, 2009.
- Office Action, U.S. Appl. No. 11/590,489 (issued), mailed Dec. 1, 2009.
- Office Action, U.S. Appl. No. 12/389,837 (pending), mailed Jun. 25, 2010.
- Office Action, U.S. Appl. No. 12/389,843 (pending), mailed Jul. 2, 2010.
- Office Action, U.S. Appl. No. 12/389,851 (pending), mailed Jul. 23, 2010.
- Office Action, U.S. Appl. No. 11/066,880 (pending), mailed Dec. 11, 2007.
- Office Action, U.S. Appl. No. 11/066,880 (pending), mailed Sep. 24, 2008.
- VoIP Watch, "All About Voice Over IP and Internet Based Telephony", Apr. 19, 2007, 2 pages.
- Kahane, et al., "Call Management Agent System Specification", VoIP-0017, Aug. 15, 1996, 36 pages.
- Sebestyen, "What is the Position of Q.2, Q.3/15 on Internet Telephony for the IMTC VoIP Forum Meeting in Seattle", ITU Telecommunication Standardization Sector, Study Group 15, Q.2&3/15 Rapporteur Meeting, Dec. 17-20, 1996, 5 pages.
- Petrack, "Backup and Alternate Gatekeepers in H.323 for High Availability, Reliability and Scalability", ITU Telecommunication Standardization Sector, Study Group 16, Ex. Q.2&3/15 Rapporteur Meeting, Feb. 1-21, 1997, 7 pages.
- Chapter 5: The CMA System, VoIP Forum, Oct. 11, 1996, 38 pages.
- Kahane, "Introduction to VocalTec's CMA System, Intelligent Switching for a New Era of Telecom", Nov. 7, 1996, 18 pages.
- CMA IPR Additions, Nov. 17, 2008, 3 pages.
- Rosenberg, et al., "Internet Telephony Gateway Location", In Proc. of IEEE INFOCOM, Mar. 1998, 9 pages.
- IMTC Voice Over IP Forum Service Interoperability Implementation Agreement, Draft 0.91, VOIP97-008, Jan. 13, 1997, 70 pages.
- "H.323: The Leading Standard in Voice Over IP", Dec. 1996, 4 pages.
- Oppen, et al., "The Clearinghouse: A Decentralized Agent for Locating Named Objects in a Distributed Environment", ACM Transactions on Office Information Systems, vol. 1, No. 3, Jul. 1983, pp. 230-235.
- Sharp, et al., "Advanced Intelligent Networks—Now a Reality", Electronics & Communication Engineering Journal, Jun. 1994, pp. 153-162.
- Freiedes, et al., "ISDN opportunities for Large Business—800 Service Customers", IEEE International Conference on Communications, Jun. 1986, vol. 1, 7 pages.
- Mark Handley, "Session Invitation Protocol," Feb. 22, 1996, 17 pages.
- Yeager, et al., "Web Server Technology, the Advanced Guide for World Wide Web Information Providers", Chapter 6: Searching for Information on the Web, 6.8.2. Using Database as Indexes to a Document Collection, Apr. 1996, 3 pages.
- Matta, et al., "Type-of-Service Routing in Dynamic Datagram Networks", Jun. 1994, Proceedings of the Conference on Computer Communications, IEEE, pp. 992-999.
- Peeren, "The Intelligent Web", IEEE Apr. 1996, vol. 1, 13 pages.
- Clark, et al., "Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanism", Oct. 1992 ACM 0-89791-526-7/92/0008/0014, pp. 14-26.
- Low, "The Internet Telephony Red Herring", Intelligent Networked Platforms Dept., Intelligent Networked Computing Laboratory, Hewlett-Packard Laboratories, Bristol, May 15, 1996, pp. 1-15.
- Waksberg, "Axe 10 and the Intelligent Network", Commutation and Transmission, No. 4, Dec. 1993, pp. 67-76.
- Harner, et al., "On the Move Project", Dec. 1996, pp. 1-44.
- Sisalem, et al., "The Network Video Terminal", IEEE, Aug. 1996 Proceedings, pp. 3-10.
- Kolarov, et al., "End-to-End Adaptive Rate Based Congestion Control Scheme for ABR Service in Wide Area ATM Networks", Feb. 1995, pp. 138-143.
- Niitsu, et al., "Computer-Aided Stepwise Service Creation Environment for Intelligent Network", Supercomm/International Conference on Communications, Jun. 14, 1992, pp. 454-458.
- Sairamesh, et al., "An Approach to Pricing, Optical Allocation and Quality of Service Provisioning in High Speed Packet Networks", Fourteenth Annual Joint conference of the IEEE Computer and Communication Societies, vol. 2, Apr. 1995, 21 pages.
- Sibal, et al., "Controlling Alternate Routing in General Mesh Packet Flow Networks", Association of Computing Machinery, Aug. 1994, pp. 168-179.
- Morgan, et al., "Service Creation Technologies for the Intelligent Network", AT&T Technical Journal, Jan. 1991, pp. 58-71.
- Fujioka, et al., "Universal Service Creation and Provision Environment for Intelligent Network", International Switching Symposium, May 1990, pp. 149-156.
- Moy, "RFC 1247—IETF Network Working Group: OSPF Version 2", Jul. 1991, <<http://tools.ietf.org/html/rfc1247>>, 183 pages.
- Ahmadi, et al., NBBS Traffic Management Overview, IBM Systems Journal, vol. 34, No. 4, Dec. 1995, pp. 640-648.
- Braden, et al., "Integrated Services in the Internet Architecture: An Overview", RFC 1633, Jul. 1994, 28 pages.
- Huitema, "Routing in the Internet", Prentice-Hall PTR, Mar. 1995, 168 pages.
- Ziegler, "A Distributed Information System Study", IBM System Journal, vol. 18, No. 3, Dec. 1979, pp. 374-401.
- Elixman, et al., "OpenSwitching-Extending Control Architectures to Facilitate Applications", International Switching Symposium, vol. 2, Apr. 23, 1995, pp. 239-243.
- Kabay, et al., "The Service Node—An Advanced IN Services Element", Technology Journal, vol. 13, No. 2, Apr. 1995, pp. 64-72.
- Maruyama, "A Concurrent Object-Oriented Switching Program in Hill", IEEE Communications Magazine, vol. 29, No. 1, Jan. 1991, pp. 60-68.
- Mayer, et al., "Service Net—2000: An Intelligent Network Evolution", AT&T Technical Journal, vol. 70, No. 3/4, Jan. 1991, pp. 99-110.
- Shabana, et al., "Intelligent Switch Architecture", Proceedings of the National Communications Forum, vol. 22, No. 2, Sep. 30, 1988, pp. 1312-1320.
- Xiao, et al., abstract, "Sharing Model of Network Node Resources and Real-Time Scheduling", Journal: Mini-Micro Systems, vol. 16, No. 12, Dec. 1995, 1 page.
- Gupta, et al., "Resource Sharing for Multi-Party Real-Time Communication", Conference Title: Proceedings IEEE INFOCOM '95, The Conference on Computer Communications, Fourteenth Annual Joint Conference of the IEEE Computer and Communications Societies, Bringing Information to people. Part vol. 3, Apr. 1995, pp. 1230-1237.
- Klein, et al., "Minimax Resource Allocation with Tree Structured Substitutable Resources", Journal: Operations Research, vol. 39, No. 2, Mar. 1991, pp. 285-295.
- Weinrib, et al., abstract, "Decentralized Resource Allocation for Distributed Systems", Conference Title: IEEE INFOCOM '87, The Conference on Computer Communications, Proceedings, Sixth Annual Conference—Global Networks: Concept to Realization, Mar. 1987, 1 page.
- Malamud, et al., "Principles of Operation for the TPC.INT Subdomain: General Principles and Policy", RFC 1530, 6 pages, Oct. 1993.
- Berners-Lee, "Universal Resource Identifiers in WWW", RFC 1630, 23 pages, Jun. 1994.
- Manning, et al., "DNS NSAP Resource Records", RFC 1706, 9 pages, Oct. 1994.
- "Web-On-Call Voice Browser Redefines Access to the Web", The Netphonic Press Release, Mar. 1996.
- Tao, et al., "Internet Access via Baseband and Broadband ISDN Gateways", IEEE International Conference on Computers and Communications, pp. 485-490 (Apr. 1994).
- "Workstation Communications System", IBM Technical Disclosure Bulletin, vol. 37, No. 39, Sep. 1994.
- Bonetti, et al., "Distribution of RFC 1327 Mapping Rules via the Internet DNS: the INFNet Distributed Gateway System", Computer Networks and ISDN Systems, vol. 27, No. 3, pp. 461-469, Dec. 1994.
- Yang, et al., "The Design and Implementation of a Service Logic Execution Environment Platform", Globecom '93, vol. 3, Nov. 1993.

US 8,204,046 B2

Page 7


- Esaki, et al., "Abstraction and Control of Transport Network Resources for Intelligent Networks", Electronics & Communications in Japan: Part 1—Communications, vol. 76, No. 1, Jan. 1993.
- Chang, et al., "Rapid Deployment of CPE-based Telecommunications Services", IEEE Globecom, vol. 2, pp. 876-880, Nov. 1994.
- Swale, et al., "Distributed Intelligence and Data in Public and Private Networks", BT Technology Journal, vol. 13, No. 2, pp. 95-104, Apr. 1995.
- Comer, et al., "Uniform Access to Internet Directory Services", Computer Communication Review, vol. 20, No. 4, pp. 50-59, Sep. 1990.
- Mach2DNS (Domain Name Services) Module Au-A-11045/97, Jun. 13, 1995.
- Muffett, "Proper Care and Feeding of Firewalls", Sun Microsystems, pp. 1-12, Nov. 14, 1994.
- Crocker, "RFC 822—Standard for the Format of ARPA Internet Text Messages", RFC 822, Internet: <<http://www.faqs.org/rfcs/rfc822.html>>, pp. 1-42, Aug. 13, 1982.
- Girod, et al., "Requirements for URN Resolution Systems", Internet: <<http://www.potaroo.net/ietf/ldref/draft-girod-urn-res-require/>>, pp. 1-31, Jun. 13, 1996.
- Hamilton, "Uniform Resource Identifiers & The Simple Discovery Protocol", Department of Computer Studies, Loughborough University of Technology, Internet: <<http://martinh.net/uris/uris.html>>, 32 pages, Jun. 20, 1995.
- Hardcastle-Kille, "RFC 1279-X.500 and Domains", RFC 1279, Internet: <<http://www.faqs.org/rfcs/rfc1279.html>>, pp. 1-12, Nov. 1991.
- Mealling, et al., "Uniform Resource Names, ISO OIDs and DNS", Internet, <<http://www.gbiv.com/protocols/uri/drafts/draft-mealling-oid-dns-00.txt>>, pp. 1-14, Nov. 22, 1995.
- Mitra, et al., "Uniform Resource Names", Internet: <<http://www.mitra.biz/uri/draft-ictf-uri-resource-names-03.txt>>, pp. 1-9, Nov. 1994.
- Rose, et al., "RFC 1486—An Experiment in Remote Printing", RFC 1486, Internet: <<http://www.faqs.org/rfcs/rfc1486.html>>, pp. 1-12, Jul. 1993.
- Shafer, et al., "URN Services", Internet: <<http://staff.oclc.org/oclc/research/publications/shafer/urn/draft-shafer-uri-urn-resolution-00.html>>, pp. 1-10, Jun. 1995.
- Shafer, et al., "URN Services", Internet: <<http://staff.oclc.org/oclc/research/publications/shafer/urn/draft-ietf-uri-urn-resolution-01.html>>, pp. 1-14, Jul. 1995.
- Sollins, et al., "RFC 1737—Functional Requirements for Uniform Resource Names", RFC 1737, Internet: <<http://www.faqs.org/rfcs/rfc1737.html>>, pp. 1-6, Dec. 1994.
- The URN Implementors, "Uniform Resource Names", D-Lib Magazine, Internet: <<http://www.dlib.org/dlib/February96/02arms.html>>, pp. 1-6, Feb. 1996.
- Everhart, et al., "RFC 1183—New DNS RR Definitions", Internet: <<http://www.faqs.org/rfcs/rfc1183.html>>, pp. 1-9, Oct. 1990.
- Mockapetris, "RFC 1101—DNS Encoding of Network Names and Other Types", Internet: <<http://www.faqs.org/rfcs/rfc1101.html>>, pp. 1-12, Apr. 1989.
- Duan, et al., "Efficient Utilization of Multiple Channels Between Two Switches in ATM IN atm Networks", IEEE Conference on Communications, Jun. 1995, pp. 1906-1911.
- Ejiri, "For Whom the Advancing Service/Network Management", IEEE Network Operations and Management Symposium, Feb. 14-17, 1994, pp. 422-433.
- Elia, et al., "Skyplex: Distribution Up-Link for Digital Television via Satellite", Apr. 1996, pp. 305-313.
- Ely, "The Service Control Point as a Cross Network Integrator", Apr. 1996, pp. 1-8.
- Feinmann, "VIC Computer Telephony", Computer Telephony, Mar. 1996, 3 pages.
- Ganor, "Talk, Talk", Tele.com, Jun. 1996, pp. 68-72.
- Grami, et al., "The Role of Satellites in the Information Superhighway", Jun. 1995, pp. 1577-1581.
- Gupta, et al., "Technical Assessment of (T)INA-TMN-OSI Technology for Service Management Applications", Feb. 1994, pp. 877-887.
- Inamori, et al., "Applying TMN to a Distributed Communications Node System with Common Platform Software", IEEE Feb. 1995, pp. 83-87.
- Jain, "Evolving Existing Narrowband Networks Broadband Networks with IN Capabilities", IEEE Apr. 1996, 6 pages.
- Kim, "Talk is Cheap", America's Network, Jul. 15, 1996, pp. 34-39.
- Kishimoto, "Agent Communication System for Multimedia Communication Service", IEEE Infocom, Mar. 1996, Proceedings vol. 1, pp. 10-17.
- Fridisch, et al., "Terminals for Accessing the Internet—The Internet Telephone", Alcatel Telecommunications Review—4th Quarter, Dec. 1996, pp. 304-309.
- Diehl, "Data's New Voice", BYTE, Sep. 1996, pp. 129-135.
- Braun, et al., "A Framework for Flow-Based Accounting on the Internet", Sep. 1993, 5 pages.
- Margulies, "CT's Cyberdate with the 'Net", Aug. 1996, Computer Telephony Periscope, pp. 28-29.
- Arango, et al., Guaranteed Internet Bandwidth, IEEE Nov. 1996, pp. 862-866.
- Perret, et al., "MAP: Mobile Assistant Programming for Large Scale Communications Networks", Apr. 1996, pp. 1128-1132.
- Platt, "Why IsoEthernet Will Change the Voice and Video Worlds", IEEE Communications Magazine, Apr. 1996, pp. 55-59.
- Retkwa, "Telephone Politics", Internet World, Jun. 1996, pp. 54-60.
- Schreyer, et al., "Least Cost Call Routing", IEEE Jun. 1996, 7 pages.
- Serrano, "Evolution of Hybrid Fibre Coaxial Network for Multimedia Interactive Services", British Telecommunications Engineering, Oct. 1996, pp. 249-253.
- Sunaga, et al., "A Reliable Communication Switching Platform for Quick Provisioning", IEEE Feb. 1995, pp. 77-82.
- Telephony on the Internet, (Workshop Information) presented by International Quality & Productivity Center, IMTC and VOICE Technology & Services News, Sep. 1996, 6 pages.
- Tsuchida, et al., "Intelligent Dynamic Service Provisioning Architecture in the Multimedia Era", Communications—Gateway to Globalization, IEEE Apr. 1996, pp. 1117-1122.
- Estrin, et al., "Design Considerations for Usage Accounting and Feedback in Internetworks", Oct. 1990, 12 pages.
- Lee, et al., "Integrated Packet Networks with Quality of Service Constraints", Global Telecommunication Conference, IEEE Dec. 1991, pp. 223-227.
- Eriksson, "MBONE: The Multicast Backbone", Communications of the ACM, vol. 57, No. 8, Aug. 1994, pp. 54-60, 13 pages.
- Gralla, "How the Internet Works", Communication on the Internet, Chapter 12, Jun. 1996, pp. 64-67.
- Kahn, "Videoconferencing Debuts on the Internet", Feb. 28, 1995, 2 pages.
- Sullivan, "Videoconferencing Arrives on the Internet", PC Week, Aug. 1996, 1 page.
- Macedonia, et al., "Mbone Provides Audio and Video Across the Internet", Apr. 1994, 13 pages.
- Schulzrinne, "RFC 1890-RTP: Profile for Audio and Video Conferencing with Minimal Control", Jan. 1996, 18 pages.
- Schulzrinne, "RFC 1889-RTP: A Transport Protocol for Real-Time Applications", Jan. 1996, 151 pages.
- Scavos, et al., "Information model: From Abstraction to Application", Feb. 1994, 13 pages.
- Sriram, et al., "Voice Packerization and Compression in Broadband ATM Networks", Apr. 1991, IEEE Journal on Selected Areas in Communications, vol. 9, No. 3, pp. 294-304.
- CT and the 'Net, "Audio and Video Over the Internet", Mar. 1996, 22 pages.
- CT and the 'Net, "Supercharging the Web with Computer Telephony", Mar. 1996, 18 pages.
- El-Gendy, H., et al., "Computer-Supported Routing for Intelligent Networks and Personalized Wired Communications", IEEE, Nov. 1994, pp. 1027-1033.
- Grigoris, "Computer Telephony Over the Internet", CT and the Net, Mar. 1996, 6 pages.
- McConnell, "How to Build an Internet PBX", <<http://www.phonezone.com/ip-phone.htm>>, Oct. 1996, pp. 1-9.
- Sears, "Innovations in Internet Telephony: The Internet as the Competitor to the POTS Network", Innovations in Internet Telephony: The Internet as the Successor to the POTS Network, Feb. 28, 1996, pp. 1-6.
- Stylus Innovation, "Computer Telephony and the Internet", <<http://www.stylus.com/hvml.htm>>, Oct. 1996, 9 pages.

US 8,204,046 B2

Page 8

- Welcome to the Phone Zone, Pacific Telephony Design, <<http://www.phonezone.com/index2.htm>>, obtained from Internet Oct. 28, 1996, pp. 1-9.
- Yang, "INETPhone: Telephone Services and Servers on Internet", RFC 1789, Apr. 1995, Oct. 1996, pp. 1-6.
- Ng, et al., "Distributed Architectures and Databases for Intelligent Personal Communication Networks", 1992 IEEE, ICWC '92, Jun. 1992, pp. 300-304.
- Malyan, et al., "Network Architecture and Signaling for Wireless Personal Communications", IEEE Journal on Selected Areas in Communications, vol. 11, No. 6, Aug. 1993, 12 pages.
- Malyan, et al., "A Microcellular Interconnection Architecture for Personal Communications Networks", 1992 IEEE, May 1992, pp. 502-505.
- Cheriton, "Dissemination-Oriented Communication Systems: Final Report", Nov. 26, 1996, 18 pages.
- Imilienski, et al., "Mobile Wireless Computing", Communications of the ACM, Oct. 1994, vol. 37, No. 10, pp. 18-28.
- Bakre, et al., "M-RPC: A Remote Procedure Call Service for Mobile Clients", 1995 ACM 0-89791-814-S/95/10, Nov. 1995, pp. 97-110.
- O'Malley, et al., "A Dynamic Network Architecture", ACM Transactions on Computer Systems, vol. 10, No. 2, May 1992, pp. 110-143.
- Cheshire, et al., "Internet Mobility 4x4", ACM 1996, <www.acm.org>, Aug. 1996, pp. 1-12.
- Chang, et al., "A Service Acquisition mechanism for the Client/Service Model in Cygnus", IBM Canada Laboratory Technical Report TR 74.059, May 1991, pp. 323-345.
- Low, et al., "Webin—An Architecture for Fast Development of IN-Based Personal Services" Workshop Record Intelligent Network. Freedom and Flexibility: Realizing the Promise of Intelligent Network Services, IEEE, Apr. 21, 1996, 6 pages.
- Johnson, "Scalable Support for Transparent Mobile Host Internetworking", Wireless Networks 1 (1995), Sep. 1995, pp. 311-321.
- Cobbold, et al., "Enhancements for Integrated Wireless Personal Communications Over Metropolitan Area Networks", IEEE, Apr. 1996, pp. 1370-1376.
- Jennings, "End-user Requirements for High-Integrity Directory", 1996, ICC '96, Conference Record, Converging Technologies for Tomorrow's Applications, 1996 IEEE International Conference on vol. 3, Jun. 23-27, 1996, pp. 1793-1796, vol. 3; Digital Object Identifier 10.1109/ICC.1996.535599.
- Anand, et al., "Accounting Architecture for Cellular Networks", 1996 IEEE, Feb. 1996, pp. 184-189.
- Civanlar, et al., "FusionNet: Joining the Internet & Phone Networks for Multimedia Applications", Nov. 1996, pp. 431-432.
- Karttunen, et al., "Cost Structure Analysis and Reference Model for SCALEABLE Network Services", The Institution of Electrical Engineers, Nov. 1996, pp. 3/1-3/9.
- Chen, et al., "ATM and Satellite Distribution of Multimedia Educational Courseware", IEEE, Jun. 1996, pp. 1133-1137.
- Aidarous, et al., "The Role of the Element Management Layer in Network Management", IEEE Network Operations and Management Symposium, Feb. 14-17, 1994, pp. 59-70.
- Pontailier, "TMN and New Network Architectures", IEEE Communications Magazine, Apr. 1993, 6 pages.
- Li, et al., "Time-Driven Priority Flow Control for Real-Time Heterogeneous Internetworking", Proceedings in Computer Communications, Fifteenth Annual Joint Conference of the IEEE Computer Communication Generation, San Francisco, Mar. 24-28, 1996, vol. 1 Conf., Mar. 24, 1996 IEEE, pp. 189-197.
- Newton, "Telephony Messaging on the Internet", Computer Telephony, Jul. 1996, pp. 10-16.
- Gareiss, "Voice Over the Internet", Data Communications, vol. 25, No. 12, Sep. 1996, pp. 93, 94, 96, 98 and 100.
- Mills, "Phone Service Via the Internet May Slash Rates", The Washington Post, Aug. 11, 1996, p. A4.
- Venditto, "Internet Phones, The Future is Calling", Internet World, Jun. 1996, pp. 40-52.
- Carrel, "The TACACS+ Protocol, Version 1.75", Network Working Group, Lol Grant, Internet-Draft Systems, Oct. 1996, draft-grant-tacacs-00.txt, 64 pages.
- Low, "The Internet Telephony Red Herring", IEEE 1996, May 15, 1996, pp. 72-80.
- Chapman, et al., "Overall Concepts and Principles of TINA Version 1.0", Telecommunications Information Networking Architecture Consortium, TINA-C Deliverable. Date of Issue: Feb. 17, 1995, 65 pages.
- Parnass, "IDT to Unveil Net2Phone At Comdex; Revolutionary System Uses Internet to Make Overseas PC-to-Telephone Calls for 10 cents a Minute", Nov. 6, 1995, <www.highbeam.com>, 2 pages.
- Fuente, et al., "Application of the TINA-C Management Architecture", Apr. 1995, 12 pages.
- Finseth, "An Access Control Protocol, Sometimes Called TACACS", University of Minnesota, Jul. 1993 Network Working Group, Request for Comments: 1492, 22 pages.
- "VocalTec's Telephony Gateway—The Ultimate Internet Telephony Solution?", Computer Telephony, Sep. 1996, p. 30, 3 pages.
- Fang, "Building an Accounting Infrastructure for the Internet", IEEE Nov. 1996, pp. 105-109.
- Lipoff, "Operations Support System Framework for Data Over Cable Services", Data Over Cable Technical Reports, TR-DOCS-OS-SIW08-961016, Oct. 16, 1996, pp. 1-20.
- Edell, et al., abstract, "Billing Users and Pricing for TCP", Journal on Selected Areas in Communications, vol. 13, Issue 7, Sep. 1995, 2 pages.
- Berners-Lee, et al., "Hypertext Transfer Protocol—HTTP/1.0", HTTP Working Group Internet Draft (Oct. 1995).
- Mockapetris, "Domain Names-Concepts and Facilities", RFC 1034, Nov. 1987.
- Mockapetris, "Domain Names-Implementation and Specification", RFC 1035, 45 pages, Nov. 1987.
- Malamud, et al., "Principles of Operation for the TPC/INT Subdomain: Remote Printing—Technical Procedures", RFC 1528, 10 pages, Oct. 1993.
- Malamud, et al., "Principles of Operation for the TPC/INT Subdomain: Remote Printing—Administrative Policies", RFC 1529, 5 pages, Oct. 1993.

* cited by examiner



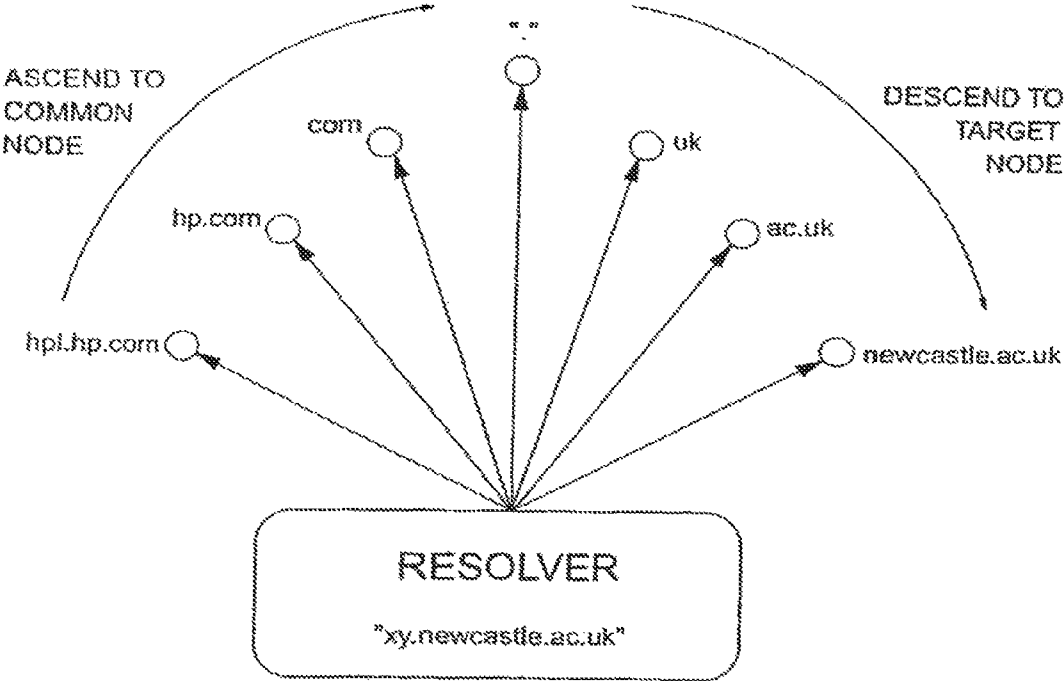



FIG. 3
(PRIOR ART)

U.S. Patent

Jun. 19, 2012

Sheet 3 of 14


US 8,204,046 B2

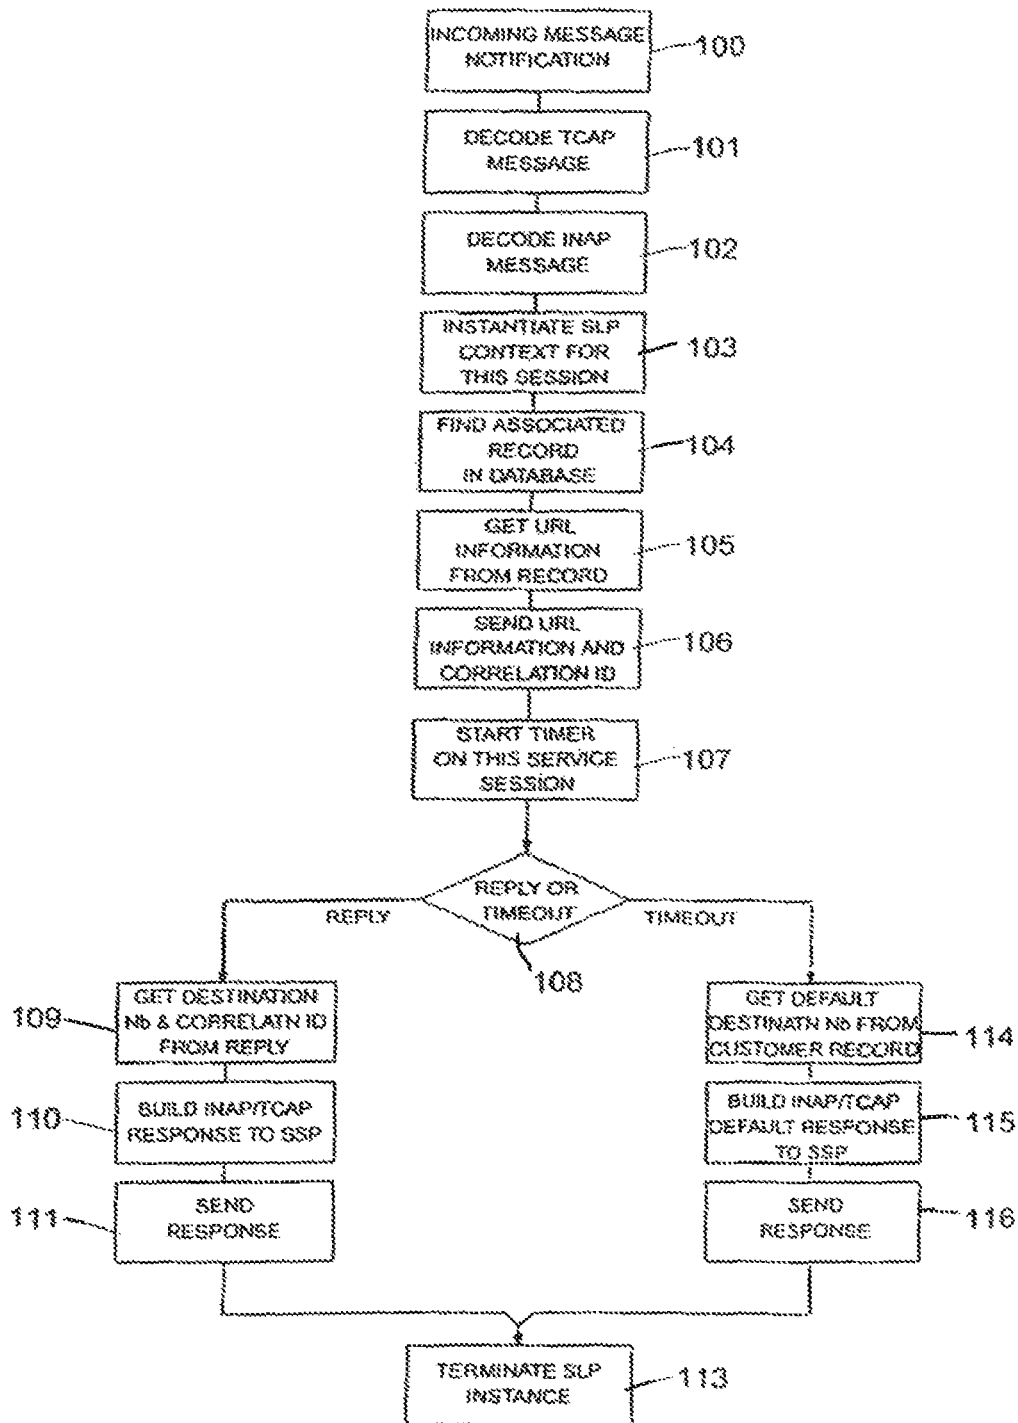


http : //www.hp.com / Products.html

SCHEME HOST LOCATION ABSOLUTE PATH

FIG. 5
(PRIOR ART)








FIG. 8

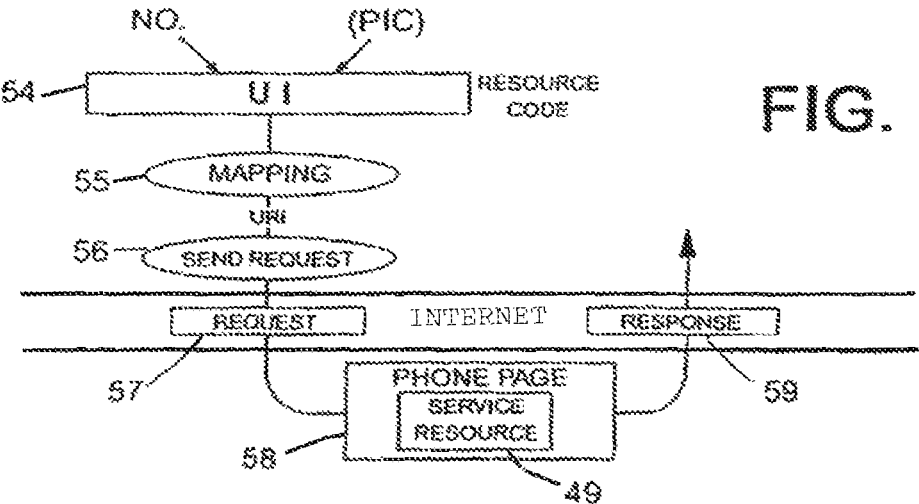


FIG. 9




FIG. 10

U.S. Patent

Jun. 19, 2012

Sheet 7 of 14

US 8,204,046 B2




FIG. 11




FIG. 12A




FIG. 12B




FIG. 12C

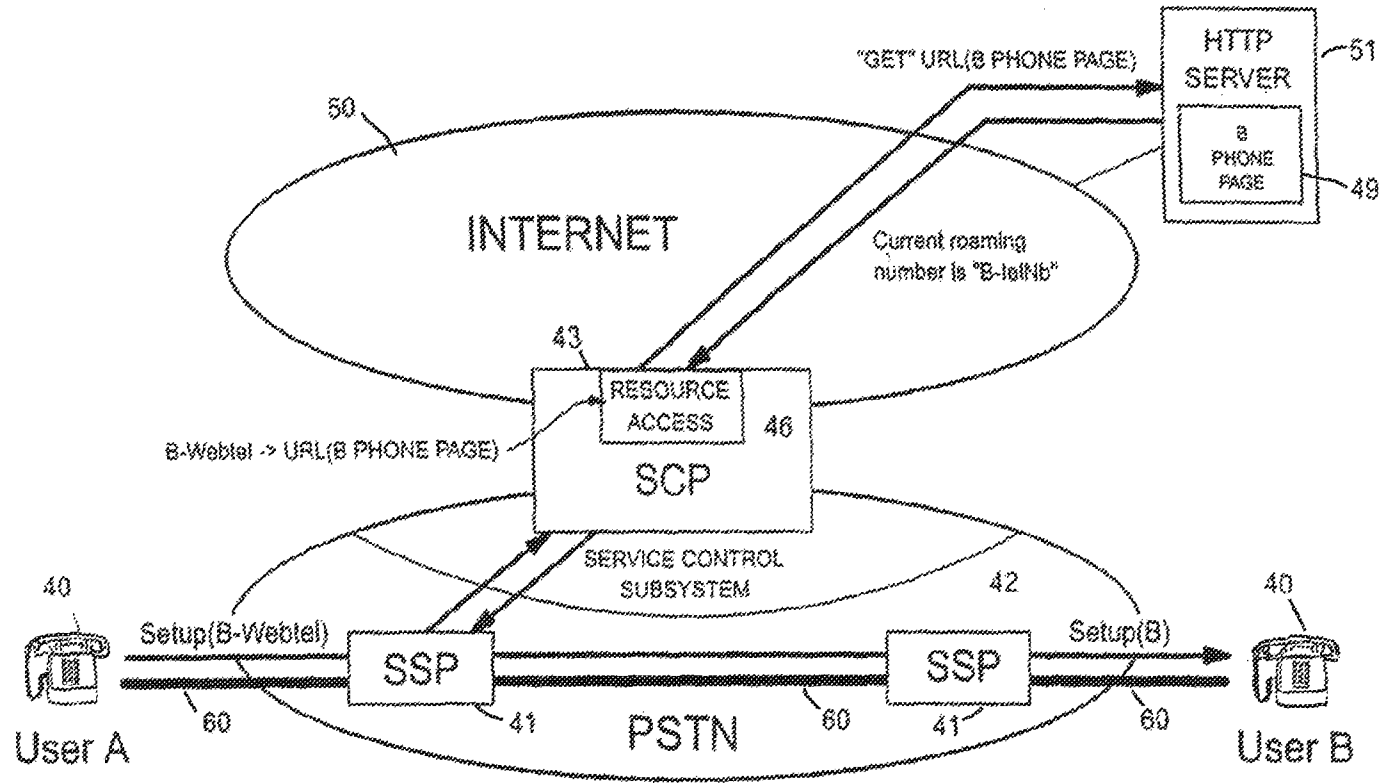


FIG. 13

A166

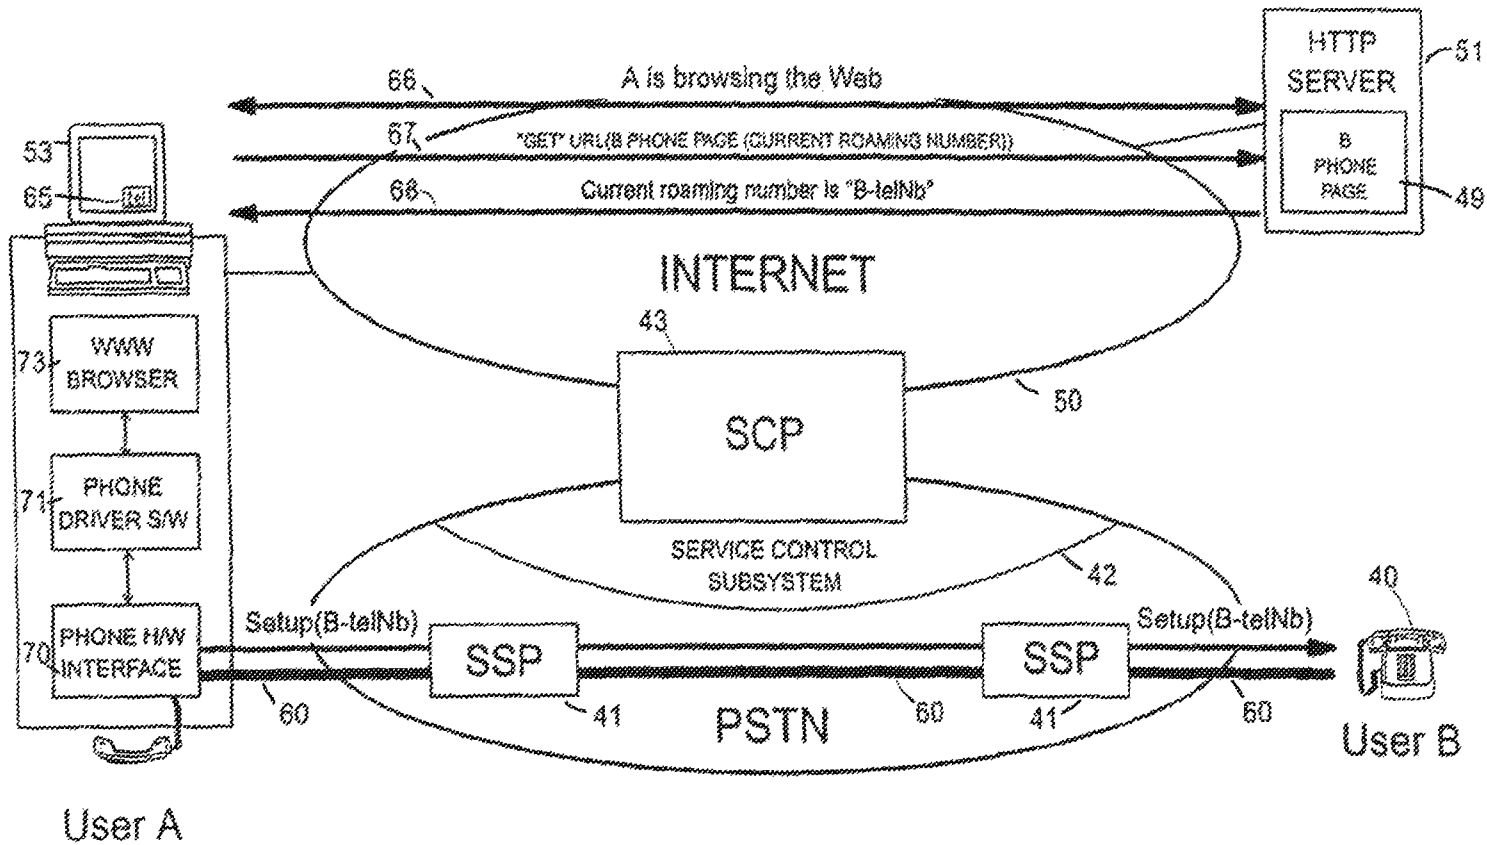


FIG. 14

A167

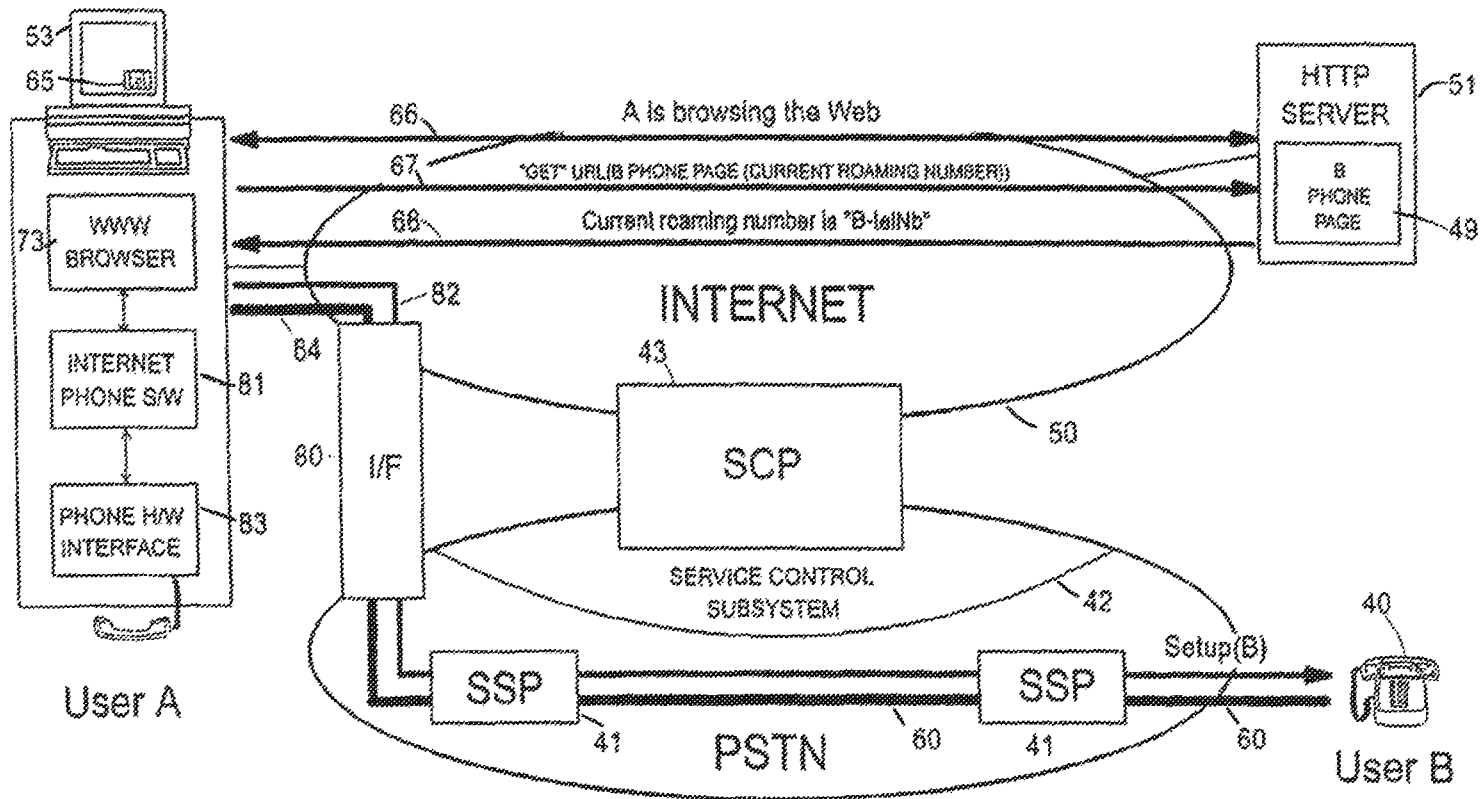


FIG. 15

A168

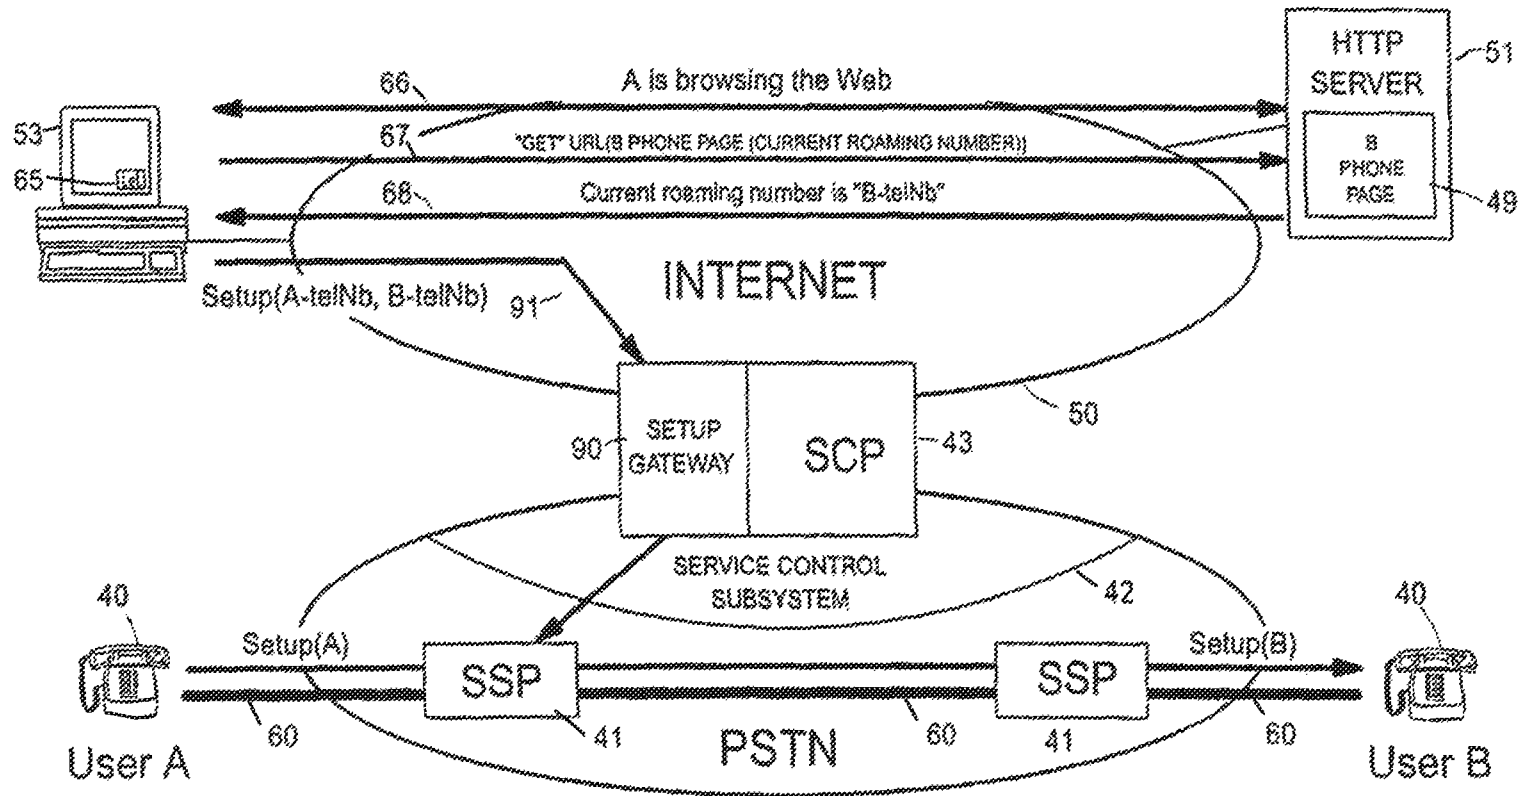


FIG. 16

A169

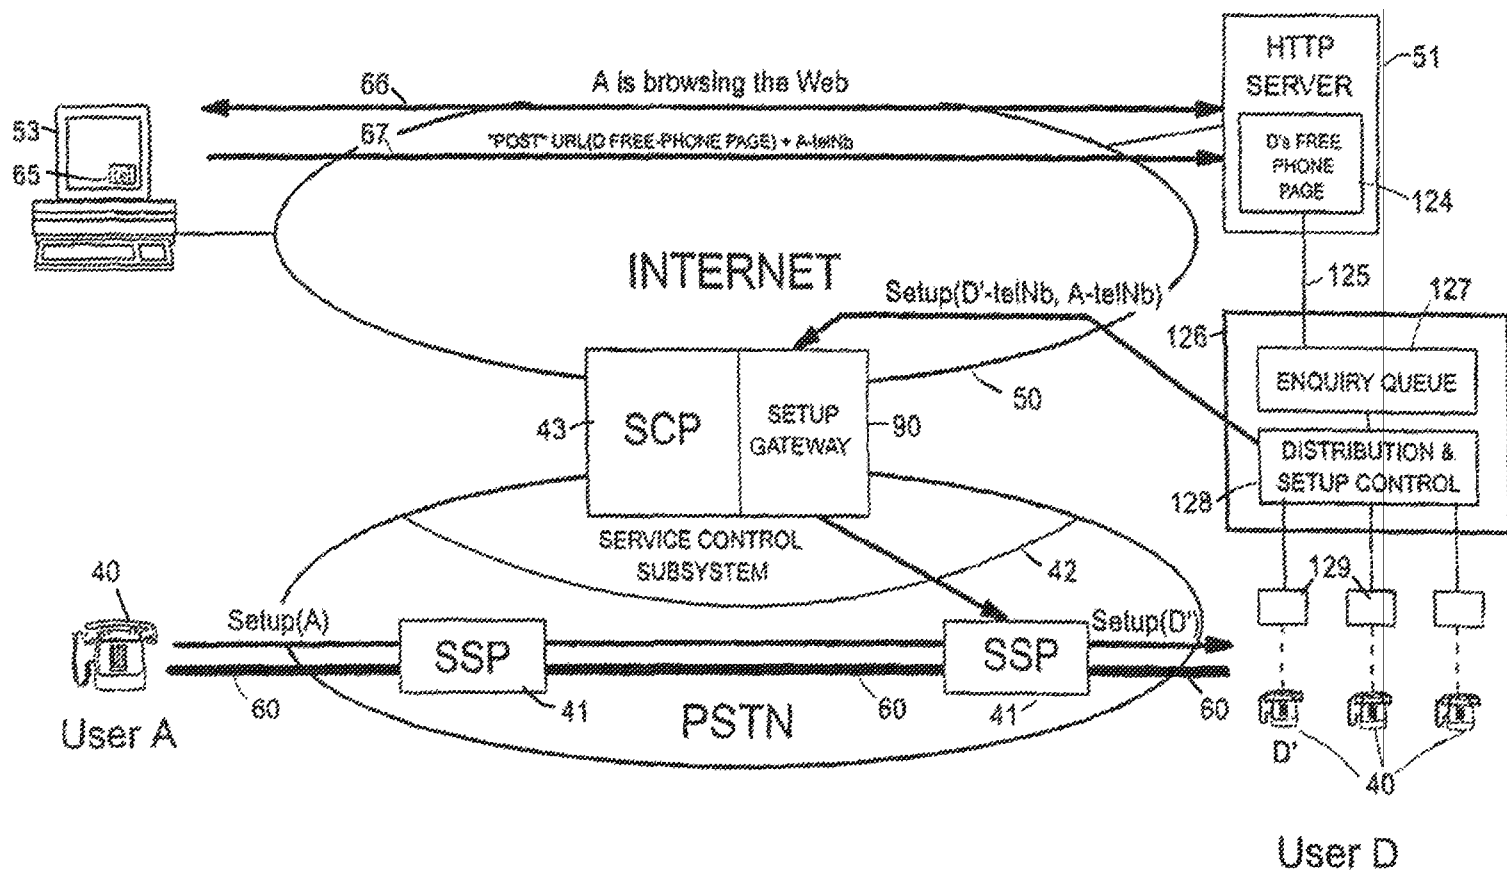


FIG. 17

A170




FIG. 18

A171

US 8,204,046 B2

1

METHOD AND APPARATUS FOR ACCESSING SERVICE RESOURCE ITEMS THAT ARE FOR USE IN A TELECOMMUNICATIONS SYSTEM

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application is a continuation of U.S. patent application Ser. No. 11/066,880, filed Feb. 25, 2005, which is a divisional of U.S. patent application Ser. No. 10/052,285, filed Jan. 18, 2002 (now U.S. Pat. No. 7,012,916), which is a divisional of U.S. patent application Ser. No. 09/077,795, filed Jun. 5, 1998 (now U.S. Pat. No. 6,466,570), which claims priority to, and is a national stage application under 35 USC 371 from, PCT application No. PCT/GB96/03055, filed Dec. 11, 1996, which further claims priority to Great Britain Patent Application No. GB 9603582.9, filed Feb. 20, 1996, to European Patent Application No. EP 95410148.1, filed Dec. 22, 1995, and to Great Britain Patent Application No. GB 9525190.6, filed Dec. 11, 1995. All of the above-listed patent applications, publications thereof, and patents, are hereby incorporated by reference herein as to their entireties.

FIELD OF THE INVENTION

The present invention relates to a method of accessing service resource items that are intended to be used in setting up bearer channels through a switched telecommunications system.

As used herein, the term “switched telecommunication system” means a system comprising a bearer network with switches for setting up a bearer channel through the network. The term “switched telecommunication system” is to be taken to include not only the existing public and private telephone systems (whether using analogue phones or ISDN-based), but also broadband (ATM) and other switch-based bearer networks that are currently being implemented or may emerge in the future. For convenience, the term “switched telecommunication system” is sometimes shortened herein to telecommunication system.

Reference to a “call” in the context of a switched telecommunication system is to be understood as meaning a communication through a bearer channel set up across the bearer network, whilst references to call setup, maintenance and takedown are to be taken to mean the processes of setting up, maintaining and taking down a bearer channel through the bearer network. Terms such as “call processing” and “call handling” are to be similarly interpreted.

The term “communication system” when used herein should be understood as having a broader meaning than switched telecommunication system, and is intended to include datagram-based communication systems where each data packet is independently routed through a bearer network without following a predetermined bearer channel.

BACKGROUND OF THE INVENTION

Telecommunication companies running PSTNs (Public Switched Telephone Networks) and PLMNs (Public Land Mobile Networks) are in the business of providing communication services and in doing so are providing increasing built-in intelligence in the form of “IN services” such as 800 number services and call forwarding. In contrast, the World Wide Web (WWW), which has seen explosive growth in recent times, is an example of an Internet-based global network providing complex information services. These two worlds, that of the large communications utilities and that of

2

the highly dynamic, pioneer-spirit WWW information culture, are uneasy companions and each plans to encroach on the domain previously occupied by the others; thus telephony services will be offered over the WWW and information services over the public communication infrastructure.

The present invention proposes technologies for a more synergetic relationship between these two worlds than is currently envisaged and in order to place the present invention in context, a review will first be given of each of these two worlds.

Telephone Networks with IN Services

The Basic PSTN. The basic service provided by a PSTN (Public Switched Telephone Network) is the interconnection of two telephones (that is, setting up a bearer channel between the telephones) according to a called-party telephone number input at the calling-party telephone. FIG. 1 is a simplified representation of a PSTN providing such a service. In particular, customer premises equipment, CPE, **10** (such as standard analogue telephones, but also more recently ISDN terminals) are connected through an access network **11** to switching points, SPs **12**. The SPs **12** form nodes in an inter-exchange network **13** made up of interconnecting trunks **14** and SPs that are controlled by control entities **15** in the SPs. The control effected by the control entities **15** is determined by signaling inputs received from the CPEs and other SPs, and involves call setup, maintenance and clearance to provide the desired bearer channel between calling CPE and called CPE. Conceptually, the PSTN may be thought of as a bearer network and a control (signaling) network, the function of the latter being to effect call control through the bearer network, namely the control of setup, maintenance and take down of bearer channels through the bearer network; in practice, the bearer and signaling networks may use the same physical circuits and even the same logical channels.

Thus, where the CPE is a traditional dumb telephone, control signaling between the CPE and its local SP is in-band signaling, that is, the signaling is carried on the same channel as used for voice; this signaling is interpreted and converted at the SPs **12** into signaling between SPs that uses a dedicated common-channel signaling network **16** (implemented nowadays using the SS7 protocol suite). Where the CPE is an ISDN terminal, signaling is carried in a separate channel directly from the CPE on an end-to-end. Modern SPs use the ISUP (ISDN User Part) SS7 protocol for inter-exchange call control signaling whether the CPE is a standard telephone or an ISDN terminal.

Telephone Numbering Plans—As certain aspects of the present invention are influenced by the structuring of telephone numbers, a brief description will now be given of the structuring of such numbers. Telephone numbers form an international, hierarchical addressing scheme based on groups of decimal digits. The top level of the hierarchy is administered by the ITU-T, which has allocated single-digit numeric codes to the major geographic zones (for example “1” for North America, “2” for Africa, “3” for Europe, “4” for Europe, “5” for South America and Cuba, etc.). Within each zone countries are assigned 2 or 3 digit codes, so that within zone 3 France is “33”, and within zone 4 the UK is “44”. Administration of the numbering plan within a country is delegated to a national body, such as the Office of Telecommunications (“OfTel”) in the UK. The following further description is based on the UK numbering plan, but the scheme described will be recognised as having widespread applicability.

In the UK all national numbers are prefixed by a code from 01 to 09 (the ‘0’ prefix is dropped in international dialing). The currently assigned codes are “01” for Geographic Area

3

Codes, "02" for Additional Geographic Area Codes, "04" for Mobile Services, "07" for Personal Numbers, and "08" for Special Service (freephone, information). Normal wireline PSTN subscriber telephone numbers are allocated from the Geographic Area Code codes, and currently only codes pre-
5 fixed by 01 are allocated. Geographic area codes are presently 3 or 4 digits (excluding the leading '0') and there are currently 638 geographic areas each with its own code. A full national UK dialed number takes two forms:

0	171	634 8700
	area code	local number (7 digit)
0	1447	456 987
	area code	local number (6 digit)

The first case has the '0' prefix, a 3 digit area code and a 7 digit local number, and the second case has the '0' prefix, a 4 digit area code, and a 6 digit local number. Further interpretation of the local number will take place within the area exchange, as even a 6 digit address space is too large for a single switch, and for a typical local area several switches may be needed to host the required number of subscriber lines. This interpretation is opaque and is a matter for the area service provider.

In the current PSTN the inherently hierarchical and geographic interpretation of telephone numbers is mirrored by the physical architecture of the network. A telephone number is structured in a way that makes it easy to route a call through the network. At each step, the prefix of the number provides information about the current routing step, and the suffix (perhaps opaquely) provides information about subsequent routing steps; as long as a switch knows how to parse a prefix and carry out a routing step, it does not need to understand the content of the suffix, which is left for subsequent routing steps. For this reason the international and national switching fabric is also organised hierarchically.

Intelligent Networks. Returning now to a consideration of the current telephone network infrastructure, in addition to basic call handling, an SP may also serve to provide what are called IN (Intelligent Network) services; in this case the SP is termed a service switching point, SSP. An SSP 25 is arranged to suspend call processing at defined points-in-call upon particular criteria being met, and to delegate the continuation of call processing to a service control subsystem providing a service control function (SCF) either in the form of a service control point, SCP 17 (see FIG. 2) or an Adjunct 18. The Adjunct 18 is directly associated with an SSP 25 whilst the SCP 17 and SSP 25 communicate with each other via an extended common channel signaling (CCS) network 16 that may include signal transfer points (STP) 19. The SCP 17 may be associated with more than one SSP 25. Both the SCP 17 and Adjunct 18 provide a service logic execution environment (SLEE) 20 in which instances of one or more service logic programs (SLP) 21 can execute. The SLEE 20 and SLP 21 together provide service control functionality for providing services to the SSP 25.

Service logic running in an SCP or Adjunct will generally make use of subscriber information stored in a service data function (SDF) 22 that may be integral with the SCP/Adjunct or partially or wholly separate therefrom. The service data function (SDF), like the service control function (SCF) forms part of the service control subsystem of the PSTN. It may be noted that some or all of the service control function may be built into the PSTN switches themselves.

4

In addition to the SCP 17 and Adjunct 18, the FIG. 2 network includes an intelligent peripheral (IP) 23. The IP 23 provides resources to the SSP 25 such as voice announcements and DTMF digit collection capabilities. The network will also include an operation system (not shown) that has a general view of the network and its services and performs functions such as network monitoring and control.

In operation, when the SSP 25 receives a call, it examines internal trigger conditions and, possibly, user information (eg dialed digits) to ascertain if the call requires a service to be provided by the service control subsystem 17, 18; the checking of trigger conditions may be carried out at several different points in call processing. Where the SSP 25 determines that a service is required it messages the service control subsystem (either SCP 17 or Adjunct 18) requesting the desired service and sending it a logic representation of the call in terms of its connectivity and call processing status. The service control subsystem then provides the requested service and this may involve either a single interaction between the SSP and service control subsystem or a session of interactions. A typical service is call forwarding which is a called-party service giving expression to an end-user requirement as simple as "if you call me on number X and it rings ten times, try calling number Y". In this case, it is the SSP local to the called end-user that triggers its associated SCP (or Adjunct) to provide this service; it will, of course, be appreciated that the SSP must be primed to know that the service is to be provided for a called number X.

The above-described model for the provision of IN services in a PSTN can also be mapped onto PLMNs (Public Land and Mobile Networks) such as GSM and other mobile networks. Control signaling in the case of a mobile subscriber is more complex because in addition to all the usual signaling requirements, there is also a need to establish where a call to a mobile subscriber should be routed; however, this is not a very different problem from a number of called-party IN services in the PSTN. Thus in GSM, the service-data function (SDF) is largely located in a system named a Home Location Register (HLR) and the service control function in a system named a Visitor Location Register (VLR) that is generally associated on a one-to-one basis with each SSP (which in GSM terminology is called a Mobile Switching Centre, MSC).

Because subscribers are mobile, the subscriber profile is transported from the HLR to whichever VLR happens to be functionally closest to be mobile subscriber, and from there the VLR operates the (fixed) service using the subscriber profile and interacts with the SSP. The HLR and VLR thus constitute a service control subsystem similar to an SCP or Adjunct with their associated databases.

It is, of course, also possible to provide IN services in private telephone systems and, in this case, the service control function and service data function are generally either integrated into a PABX (Private Automatic Branch Exchange) or provided by a local computer. The service control subsystem, whilst present, may thus not be a physically distinct from the PABX.

The above-described general architectural framework for providing IN services has both strengths and flaws. Its main strength is that it works and many services have been successfully deployed, such as 800 number services, credit card calling, voicemail, and various call waiting and redirection services. However, despite years of standardisation, services are still implemented one-at-a-time on proprietary platforms and do not scale well. The approach has been based on large, fault-tolerant systems which provide services for hundreds of thousands or even millions of subscribers and take years to

US 8,204,046 B2

5

deploy. Furthermore, since the networks used to support these services also constitute the basic telephone infrastructure, anything attached to these networks must be rigorously vetted. Additionally, each country and operator tends to have local variations of the so-called standards making it difficult to supply standard products and thereby braking the dynamics of competition.

The World Wide Web

In contrast to the slow deliberate progress of the telephone infrastructure, the WWW has grown explosively from its inception in 1989 to become the primary electronic information distribution service in terms of spread, availability and richness of information content. Anyone can, for a modest outlay, become an information provider with a world-wide audience in a highly interconnected information architecture.

The WWW is a client-server application running over the Internet and using a client-server protocol which mandates only the simplest of exchanges between client and server. This protocol is HTTP (Hyper Text Transfer Protocol) which is optimised for use over TCP/IP networks such as the Internet; the HTTP protocol can, however, also be used over networks using different communication protocol stacks.

Since the availability of literature concerning the WWW has seen the same sort of growth as the WWW itself, a detailed description of the WWW, HTTP and the Internet will not be given herein. An outline description will, however, be given with attention being paid to certain features of relevance to the present invention.

The WWW uses the Internet for interconnectivity. Internet is a system that connects together networks on a worldwide basis. Internet is based on the TCP/IP protocol suite and provides connectivity to networks that also use TCP/IP. For an entity to have a presence on the Internet, it needs both access to a network connected to the Internet and an IP address. IP addresses are hierarchically structured. Generally an entity will be identified at the user level by a name that can be resolved into the corresponding IP address by the Domain Name System (DNS) of the Internet. Because the DNS or adaptations of it are fundamental to at least certain embodiments of the invention described hereinafter, a description will next be given of the general form and operation of the DNS.

The Domain Name System—The DNS is a global, distributed, database, and without its performance, resilience and scalability much of the Internet would not exist in its current form. The DNS, in response to a client request, serves to associate an Internet host domain name with one or more Registration Records (RR) of differing types, the most common being an address (A) record (such as 15.144.8.69) and mail exchanger (MX) records (used to identify a domain host configured to accept electronic mail for a domain). The RRs are distributed across DNS name servers world-wide, these servers cooperating to provide the domain name translation service; no single DNS server contains more than a small part of the global database, but each server knows how to locate DNS servers which are “closer” to the data than it is. For present purposes, the main characteristics of the DNS of interest are:

The host name space is organised as a tree-structured hierarchy of nodes with each host having a corresponding leaf node; each node has a label (except the root node) and each label begins with an alphabetic character and is followed by a sequence of alphabetic characters or digits. The full, or “fully qualified” name of a host is the string of node labels, each separated by a “.”, from the corresponding leaf node to the root node of the hierarchy, this latter being represented by a terminating “.” in the name. Thus a host machine “fred” of

6

Hewlett-Packard Laboratories in Bristol, England will have a fully qualified domain name of “fred.hpl.hp.com.” (note that if a host name does not have a terminal “.” it is interpreted relative to the current node of the naming hierarchy).

Each host has one or more associated Registration Records (RRs).

There are a plurality of DNS servers each with responsibility for a subtree of the name space. A DNS server will hold RRs for all or part of its subtree—in the latter case it delegates responsibility for the remainder of the subtree to one or more further DNS servers. A DNS server knows the address of any server to which it has delegated responsibility and also the address of the server which has given it the responsibility for the subtree it manages. The DNS servers thus point to each other in a structuring reflecting that of the naming hierarchy.

An application wishing to make use of the DNS does so through an associated “resolver” that knows the address of at least one DNS server. When a DNS server is asked by this resolver for an RR of a specified host, it will return either the requested RR or the address of a DNS server closer to the server holding the RR in terms of traversal of the naming hierarchy. In effect, the hierarchy of the servers is ascended until a server is reached that also has responsibility for the domain name to be resolved; thereafter, the DNS server hierarchy is descended down to the server holding the RR for the domain name to be resolved.

The DNS uses a predetermined message format (in fact, it is the same for query and response) and uses the IP protocols.

These characteristics of the DNS may be considered as defining a “DNS-type” system always allowing for minor variations such as in label syntax, how the labels are combined (ordering, separators), the message format details, evolutions of the IP protocols etc.

Due to the hierarchical naming structure, it is possible to delegate responsibility for administering domains (subtrees) of the name space recursively. Thus, the top-level domains are administered by InterNic (these top-level domains include the familiar ‘com’, ‘edu’, ‘org’, ‘int’, ‘net’, ‘mil’ domains as well as top-level country domains specified by standard two-letter codes such as ‘us’, ‘uk’, ‘fr’ etc.). At the next level, by way of example Hewlett-Packard Company is responsible for all names ending in ‘hp.com’ and British Universities are collectively responsible for all names ending in ‘ac.uk’. Descending further, and again by way of example, administration of the domain ‘hpl.hp.com’ is the responsibility of Hewlett-Packard Laboratories and administration of the subtree (domain) ‘newcastle.ac.uk’ is the responsibility of the University of Newcastle-upon-Tyne.

FIG. 3 illustrates the progress of an example query made from within Hewlett-Packard Laboratories. The host domain name to be resolved is ‘xy.newcastle.ac.uk’, a hypothetical machine at the University of Newcastle, United Kingdom. The query is presented to the DNS server responsible for the “hpl.hp.com” subtree. This server does not hold the requested RR and so responds with the address of the “hp.com” DNS server; this server is then queried and responds with the address of the ‘com’ DNS server which in turn responds with the address of the ‘.’ (root) DNS server. The query then proceeds iteratively down the ‘uk’ branch until the ‘newcastle.ac.uk’ server responds with the RR record for the name ‘xy’ in its subtree.

This looks extremely inefficient, but DNS servers are designed to build a dynamic cache, and are initialised with the addresses of several root servers, so in practice most of the iterative queries never take place. In this case the ‘hpl.hp.com’ DNS server will know the addresses of several root servers, and will likely have the addresses of ‘uk’ and ‘ac.uk’

servers in its cache. The first query to the ‘hpl.hp.com’ server will return the address of the ‘ac.uk’ server. The second query to the ‘ac.uk’ server will return the address of the ‘newcastle.ac.uk’ server, and the third query will return the RR in question. Any future queries with a ‘newcastle.ac.uk’ prefix will go direct to the newcastle DNS server as that address will be retained in the “hpl.hp.com” DNS server cache. In practice names within a local subtree are resolved in a single query, and names outside the local subtree are resolved in two or three queries.

Rather than a resolver being responsible for carrying out the series of query iterations required to resolve a domain name, the resolver may specify its first query to be recursive in which case the receiving DNS server is responsible for resolving the query (if it cannot directly return the requested RR, it will itself issue a recursive query to a ‘closer’ DNS server, and so on).

It should also be noted that in practice each DNS server will be replicated, that is, organised as a primary and one or more secondaries. A primary DNS nameserver initialises itself from a database maintained on a local file system, while a secondary initialises itself by transferring information from a primary. A subtree will normally have one primary nameserver and anything up to ten secondaries—the limitation tends to be the time required by the secondaries to update their databases from the primary. The primary database is the master source of subtree information and is maintained by the domain DNS administrator. The secondaries are not simply standby secondaries but each actively participates in the DNS with dependent servers that point to it rather than to the corresponding primary.

DNS implementations, such as BIND, are widely available as a standard part of most UNIX systems, and can claim to be among the most robust and widely used distributed applications in existence.

Operation of the WWW Referring now to FIG. 4 of the accompanying drawings, access to the Internet 30 may be by direct connection to a network that is itself directly or indirectly connected to the Internet; such an arrangement is represented by terminal 31 in FIG. 4 (this terminal may, for example, be a Unix workstation or a PC). Having a connection to the Internet of this form is known as having ‘network access’. Any entity that has network access to the Internet may act as a server on the Internet provided it has sufficient associated functionality; in FIG. 4, entity 32 with file store 37 acts as a server.

Many users of the WWW do not have network access to the Internet but instead access the Internet via an Internet service provider, ISP, 33 that does have network access. In this case, the user terminal 34 will generally communicate with the ISP 33 over the public telephone system using a modem and employing either SLIP (Serial Line Interface Protocol) or PPP (Point-to-Point Protocol). These protocols allow Internet packets to traverse ordinary telephone lines. Access to the Internet of this form is known as “dialup IP” access. With this access method, the user terminal 34 is temporarily allocated an IP address during each user session; however, since this IP address may differ between sessions, it is not practical for the entity 34 to act as a server.

A cornerstone of the WWW is its ability to address particular information resources by means of an Uniform Resource Identifier (URI) that will generally be either a Uniform Resource Locator (URL) that identifies a resource by location, or a Uniform Resource Name (URN) that can be resolved into an URL. By way of example, a full or “absolute” URL will comprise the following elements:

scheme	this is the access scheme to be used to access the resource of interest;
host	the Internet host domain name or IP address;
port	the host port for the (TCP) connection;
abs-path	the absolute path of the resource on the host.

In fact, the ‘port’ may be omitted in which case port 80 is assumed.

FIG. 5 of the accompanying drawings shows an example URL for the Hewlett-Packard products welcome page. In this case, the elements are:

scheme	http
host	www.hp.com
port	omitted (port 80 assumed)
abs-path	Products.html

The HTTP protocol is based on a request/response paradigm. Referring again to FIG. 4 of the drawings, given a particular URI identifying a resource 30 to be accessed, a client establishes a connection with the server 31 corresponding to the “host” element of the URI and sends a request to the server. This request includes a request method, and the “Request-URI” (which is generally just the absolute path of the resource on the server as identified by the “abs-path” element of the URI); the request may include additional data elements. The server 31 then accesses the resource 36 (here held on storage 37) and responds and this response may include an entity of a type identified by a MIME (Multipurpose Internet Mail Extensions) type also included in the response.

The two main request methods are:

GET—This method results in the retrieval of whatever information (in the form of an entity) is identified by the Request-URI. It is important to note that if the Request-URI refers to a data-producing process, it is the produced data which is returned as the entity in the response and not the source text of the process.

POST—This method is used to request that the destination server accept the entity enclosed in the request as a new subordinate of the resource identified by the Request-URI. The POST method can be used for annotation of existing resources, providing a message to a bulletin board, providing data to a data-handling process (for example, data produced as the result of submitting a form), and extending a database through an append operation.

In summary, the GET method can be used to directly retrieve data, or to trigger any process that will return an entity (which may either be data or a simply an indication of the result of running the process). The POST method is used for registering data and specifying this method is also effective to trigger a process in the server to handle the posted data appropriately.

The passing of information to a process triggered to run on a server using either the GET or POST method is currently done according to an interface called the Common Gateway Interface (CGI). The receiving process is often written in a scripting language though this is not essential. Typically, the triggered server script is used for interfacing to a database to service a query included in a GET request. Another use, already referred to, is to append data associated with a POST request to a database.

Other important factors in the success of the WWW is the use of the HyperText Markup Language (HTML) for repre-

US 8,204,046 B2

9

senting the makeup of documents transferred over the Www, and the availability of powerful graphical Web browsers, such as Netscape and Mosaic, for interpreting such documents in a client terminal to present them to a user. Basically, HTML is used to identify each part of a document, such as a title, or a graphic, and it is then up to the browser running in the client terminal to decide how to display each document part. However, HTML is more than this—it also enables a URI and a request method to be associated with any element of a document (such as a particular word or an image) so that when a user points to and clicks on that element, the resource identified by the URI is accessed according to the scheme (protocol) and request method specified. This arrangement provides a hyperlink from one document to another. Using such hyperlinks, a user at a client terminal can skip effortlessly from one document downloaded from a server on one side of the world, to another document located on a server on the other side of the world. Since a document created by one author may include a hyperlink to a document created by another, an extremely powerful document cross-referring system results with no central bureaucratic control.

Hyperlinks are not the only intelligence that can be built into an HTML document. Another powerful feature is the ability to fill in a downloaded “Form” document on screen and then activate a ‘commit’ graphical button in order to have the entered information passed to a resource (such as a database) designed to collect such information. This is achieved by associating the POST request method with the ‘commit’ button together with the URI of the database resource; activating the ‘commit’ button results in the entered information being posted to the identified resource where it is appropriately handled.

Another powerful possibility is the association of program code (generally scripts to be interpreted) with particular documents elements, such as graphical buttons, this code being executed upon the button being activated. This opens up the possibility of users downloading program code from a resource and then running the code.

It will be appreciated by persons skilled in the art that HTML is only one of several currently available scripting languages delivering the functionality outlined above and it may be expected that any serious Web browser will have built-in support for multiple scripting languages. For example, Netscape 2.0 supports HTML 3.0, Java and LiveScript (the latter being Netscape proprietary scripting Language).

The importance of the role of the graphical Web browser itself should not be overlooked. As well as the ability to support multiple scripting languages, a Web browser should provide built-in support for standard media types, and the ability to load and execute programs in the client, amongst other features. These browsers may be viewed as operating systems for WWW interaction.

WWW and the Telephone Network

It is possible to provide a telephony service over the Internet between connected terminals by digitising voice input and sending it over the Internet in discrete packets for reassembly at the receiving terminal. This is an example of a communication service on the Internet. Conversely, it is possible to point to a variety of information services provided over the telephone system, such as the Minitel system widely available in France. However, these encroachments into each others traditional territories pose no real threat to either the Internet or the public telephone system.

Of more interest are areas of cooperative use of the Internet and the telephone system. In fact, one such area has existed for some considerable time and has been outlined above with

10

reference to FIG. 4, namely the use of a modem link over the PSTN from a user computer 34 to an Internet service provider 33 in order to obtain dialup IP access to the Internet. This cooperative use is of a very simple nature, namely the setting up of a bearer channel over the PSTN for subsequently generated Internet traffic; there is no true interaction between the Internet and the PSTN.

Another known example of the cooperative use of the Internet and PSTN is a recently launched service by which an Internet user with a sound card in his/her terminal computer can make a voice call to a standard telephone anywhere in the world. This is achieved by transferring digitised voice over the Internet to a service provider near the destination telephone; this service provider then connects into the local PSTN to access the desired phone and transfers across into the local PSTN the voice traffic received over the Internet. Voice input from the called telephone is handled in the reverse manner. Key to this service is the ability to identify the service provider local (in telephony charging terms) to the destination phone. This arrangement, whilst offering the prospect of competition for the telecom operators for long distance calls, is again a simple chaining together of the Internet and PSTN. It may, however, be noted that in this case it is necessary to provide at least a minimum of feedback to the Internet calling party on the progress of call set to the destination telephone over the PSTN local to that telephone; this feedback need only be in terms of whether or not the call has succeeded.

From the foregoing it can be seen that the current cooperative use of the Internet and telephone system is at a very simple level.

It is an object of the present invention to provide a method of accessing a service resource item over a communications network that facilitates the integration of the PSTN and the WWW.

SUMMARY OF THE INVENTION

According to the present invention, there is provided a method of accessing service resource items for use in respect of setting up bearer channels through a switched telecommunications system, the method including the steps of:

(a)—provisioning at least one server connected to a computer network with a plurality of service resource items that are thereafter locatable on said computer network by corresponding known URIs, said computer network being logically distinct from the telecommunications system, and said service resource items relating to setup control for bearer channels through said telecommunications system with each said service resource item being associated with a respective predetermined code, said predetermined codes being distinct from said URIs and identifying end-point entities for said bearer channels;

(b)—providing a mapping between each said predetermined code and the said known URI of the service resource item associated with that predetermined code; and

(c)—utilising a said predetermined code to access a corresponding said service resource item by using said mapping to determine the URI corresponding to that resource item and then using this URI to access the service resource item over said computer network.

In one embodiment, at least some of the URIs are derivable from their corresponding said predetermined codes by manipulation according to a function specified by said mapping. In another embodiment, at least some of the URIs are derivable from their corresponding said predetermined codes by look up in an association table associating said predetermined codes and URIs according to said mapping. This asso-

US 8,204,046 B2

11

ciation table can advantageously be held on at least one database server connected to the computer network, step (c) involving accessing the database server over the computer network to determine the URI corresponding to the said predetermined code. Preferably, the said at least one database server is provided by a DNS-type distributed database system in which the URIs are held in records associated with respective names, herein referred to as domain names, by which the records can be retrieved. In this case, step (c) involves translating said predetermined code into a corresponding domain name and using this domain name to retrieve the URI of the required service resource item from the DNS-type distributed database system.

More than one service resource item can be located at the same URI; in this case, the predetermined codes of these service resource items will include respective relative-resource-identifier values that can be used at the server holding the service resource items to identify the required resource item amongst the service resource items at the same URI.

The telecommunications system may be a telephone system with each said predetermined code being either the telephone number of the calling party or the telephone number of the called party (these numbers may either be the numbers of specific telephones, or personal numbers). In one preferred embodiment where at least some of said predetermined codes are called-party telephone numbers, the corresponding service resource items are the current telephone numbers of the called parties.

Generally as regards the nature of the service resources, these may be of the following type:

service logic intended to be executed by the corresponding server upon being accessed with the result of this execution being returned to the accessing entity;

downloadable service data which upon being accessed is intended to be downloaded to the accessing entity;

downloadable service logic which upon being accessed is intended to be downloaded to the accessing entity for execution thereby.

Preferably, where URIs are referred to in the foregoing, these URIs are URLs and/or URNs. Furthermore, the servers referred to are preferably HTTP servers.

It is to be understood that reference in the foregoing to the computer network being logically distinct from the telecommunications system is not to be taken to imply that there is physical separation of the two—indeed, there will frequently be joint use of the same physical infrastructure. Furthermore, not only may bearer channels set up in the telecommunications system share the same transmission medium as the computer network, but such a bearer channel may act as a pipe for traffic across the computer network. The intention of requiring the computer network to be logically distinct from the telecommunications system is to exclude computer networks that are dedicated to the management or monitoring of the bearer network and effectively form part of the telecommunications system itself.

Preferably, the computer network is generally accessible to users of the telecommunications system as this provides a number of benefits to users that will become apparent hereinafter. The phrase “generally accessible” should not be construed as meaning that all users of the telecommunications system have such access to the computer network or can get such access but, rather, it should be understood as meaning that a significant proportion of these users have or can obtain access to the computer network.

By way of example, in one preferred embodiment of the invention, the computer network generally accessible to users of the telecommunications system but logically distinct from

12

it, is the Internet and the telecommunications system is a public telephone system. In another embodiment, the telecommunication system is a private system including a PABX, and the computer network is a LAN.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will now be described, by way of non-limiting example, with reference to the accompanying diagrammatic drawings, in which:

FIG. 1 is a simplified diagram of a standard PSTN;

FIG. 2 is a simplified diagram of a known PSTN with IN service capability;

FIG. 3 is a diagram illustrating host domain name resolution by the DNS of the Internet;

FIG. 4 is a diagram illustrating the functioning of the World Wide Web;

FIG. 5 is a diagram illustrating the format of a standard URL;

FIG. 6 is a diagram of a first arrangement in which service resource items are held on HTTP servers accessible both by the service control subsystem of a PSTN and by Web users;

FIG. 7 is a diagram illustrating the processing of a service request by the SCP of FIG. 6;

FIG. 8 is a diagram illustrating the format of a resource code used by the FIG. 6 SCP when accessing a service resource item;

FIG. 9 is a diagram illustrating the process of accessing a service resource in the case where the service code does not include an RRI part;

FIG. 10 is a diagram illustrating the process of accessing a service resource in the case where the service code includes an RRI part;

FIG. 11 is a diagram illustrating the derivation of the URI of a service resource by parsing an input telephone number in accordance with the present invention;

FIG. 12A is a diagram depicting a name space (the “telname space”) constituted by the domain names derived by a parsing of a predetermined set of telephone numbers;

FIG. 12B is a diagram depicting the incorporation of the telname space without fragmentation into the DNS;

FIG. 12C is a diagram depicting the incorporation of the telname space in fragmented form into the DNS;

FIG. 13 is a diagram illustrating the overall operation of the FIG. 6 arrangement in providing a roaming number service in response to a telephone number being dialed at a standard phone;

FIG. 14 is a diagram illustrating the overall operation of the FIG. 6 arrangement when utilised by a Web user in setting up a call through a telephone interface integrated into the user’s Web terminal;

FIG. 15 is a diagram illustrating the overall operation of an arrangement in which an interface is provided between the PSTN and the Internet for telephone traffic;

FIG. 16 is a diagram illustrating the overall operation of an arrangement in which a call setup gateway is provided between the Internet and the PSTN;

FIG. 17 is a diagram illustrating the overall operation of an arrangement in which a freephone service is implemented for Web users; and

FIG. 18 is a diagram similar to FIG. 6 illustrating the provision of a distributed processing environment for inter-connecting elements of the service control subsystem of the PSTN.

BEST MODE OF CARRYING OUT THE INVENTION

FIG. 6 illustrates an arrangement for the provision of services in a PSTN conventionally comprising an inter-ex-

13

change network 13 (including trunks and switches at least some of which are SSPs 41 with associated IPs), an access network 11 connecting customer premise equipment (here shown as telephones 40) to the network 13, and a service control subsystem 42 including at least one SCP for providing services to the SSPs 41 upon request. It will be appreciated that the FIG. 6 representation of a PSTN is highly diagrammatic.

The SCP 43 may operate in a conventional manner responding to service requests from SSPs 41 to run specific service logic on particular data according to information contained in the service request, and to send back to the requesting SSP appropriate instructions for effecting call set up. A service request is generated by the SSP in response to predetermined trigger conditions being met at a trigger check point, there being one or more such check points in the course of handling a call (it may be noted that where the trigger conditions have been downloaded to the SSP from the SCP then it could be said that the SSP is responding to an information request by the SCP when contacting the SCP upon the trigger conditions being met—however, in the present specification, this initial communication from the SSP to the SCP will be referred to as a “service request”).

The SCP 43 is also provided with a network access interface 44 to the Internet 50 in order to make use of certain service resource items 49 (also referred to below simply as “service resources”) during the course of processing at least certain service requests from the SSPs 41. These service resources 49 are held as WWW pages on HTTP servers 51 (more particularly, on service resource databases 52 of these servers 51). The WWW pages containing these service resources are referred to below as “phone” pages. The servers 51 are connected to the Internet and the phone pages are read accessible using respective URLs or URNs (for convenience, the more general term URI will be used hereinafter to mean the Internet-resolvable indicator of the location of a phone page).

The service resources may be service logic or service data and may be used by an otherwise standard service logic program running on the SCP, by accessing the phone page of the required resource using the appropriate URI. In certain cases, the service resources 49 may provide substantially all of the service control and data associated with a particular service. In this case, the service logic program running in the SCP 43 is of skeleton form, being instantiated on receipt of a service request and then serving to initiate service resources access and to return the results of this access to the entity that made the service request. In fact, according to this approach, the SCP could be implemented simply as a platform for fetching and executing phone-page service logic and would not need to have the complex provisioning and management systems for such logic as is required by standard SCP platforms; SCPs could then become more ubiquitous, possibly being associated with every SSP.

FIG. 7 is a flow chart illustrating the progress of events in the case where the SCP 43 handles a service request by accessing a phone-page service resource. Upon receipt of a service request in an INAP message (step 100), SCP 43 decodes the TCAP/INAP message structure in standard manner (steps 101 and 102) well understood by persons skilled in the art. Next, SCP 43 instantiates a service logic program, SLP, to handle the request (step 103). This SLP is then responsible for looking up the URL of the required service resource as determined from information contained in the service request (steps 104, 105). For example, if the service request relates to a called-party service, then the required resource will be indicated by the dialed number and the latter will be

14

used to derive the URL of the resource. Once the URL of the desired service resource has been ascertained, a resource request (for example, in the form of an HTTP request message) is sent over the Internet to the corresponding server holding the desired service resource (step 106); a correlation ID is also passed with the resource request to enable a response from the latter to be linked with the appropriate SLP instance. A timer is also started (step 107).

If a response is received from the accessed resource before the expiration of a time-out period (tested in step 108), then the response, which is usually in the form of a destination number, is supplied to the appropriate SLP as identified using the correlation ID passed with the response (step 109). An INAP/TCAP response message is then prepared and sent to the entity that made the original service request (steps 110 and 111) after which the SLP instance is terminated (113).

If in step 108, a time-out occurs before a response is received, then a default response value (generally a default destination number) may be looked up in the customer record and put in an INAP/TCAP message and sent back to the requesting entity (steps 114 to 116). The SLP instance is then terminated (113).

Locating & Accessing Service Resources

The functionality associated with accessing a phone-page resource is schematically represented in FIG. 6 by resource access block 46. Block 46 includes URI determination block 47 for determining the URI of the phone page containing the desired resource on the basis of parameters passed to block 46. Using the URI returned by block 47, the resource access block 46 then accesses the phone page of the required service resource 49 over the Internet through interface 44.

Resource Codes—It is possible that more than one service resource is associated with a particular telephone number; in this case the resource access block 46 will need to know additional information (such as current point-in-call, pic) to enable the appropriate service resource to be identified. If the service resources associated with a number are located on different phone pages, then the additional information is also passed to the URI determination block 47 to enable it to return the URI of the appropriate phone page. It is also possible for all the service resources associated with a number to be located on the same phone page. In this case, the resource access block 46 uses the additional information to pass a resource-identifying parameter with its access request to the phone page concerned; it is then up to the functionality associated with the phone page to access the correct service resource.

Thus, each service resource can be considered as being identified by a respective resource code 54 (see FIG. 8) made up of a first part UI (“URI Identifier”) used to identify the URI at which the resource is located on the Internet, and a second part RRI (“Relative Resource Identifier”) used to identify the resource amongst plural resources at the same URI.

Resource Access—Where only one service resource 49 is located on a phone page 58 identified by a unique URI, then the resource code 54 simply comprises the UI, generally either a telephone number alone or a telephone number plus a pic parameter (see FIG. 9). In this case, accessing a resource simply involves mapping the whole resource code 54 into the corresponding URI (process 55) and then sending a request 57 to the corresponding phone page 58, this latter itself constituting the desired service resource 49. The result of accessing resource 49 is then returned in response message 59.

In contrast, where multiple service resources 49 are located on the same phone page 58 (FIG. 10), the resource code 54 comprises both a UI and RRI, the UI generally being a telephone number and the RRI a pic or other parameter for

US 8,204,046 B2

15

distinguishing between the co-located resources. In this case, accessing a resource involves mapping the UI part of the resource code **54** into the corresponding URI (process **55**) and then sending a request **57** to the corresponding phone page (process **56**), the request including the RRI of the resource code. The phone page **58** includes functionality **64** for accessing the required resource on the basis of the RRI in the request message. The result of accessing the required resource **49** is then returned in response message **59**.

An alternative to the FIG. **10** method of accessing a service resource that is co-located with other resources on a phone page, would be to retrieve the whole page across the Internet simply using the URI derived from the UI part of the resource code, and then to extract the desired resource on the basis of the RRI.

URI Determination from Resource Code—The implementation of the URI determination block **47** that performs process **55** will next be considered. Block **47** may be implemented in a variety of ways, four of which are described below:

Direct Input

It would be possible, though not necessarily convenient, to arrange for the calling party to input directly the required URI. The calling party may thus input the host id component of the URI required (either in the form of a host domain name or host IP address) plus the path component of the URI. For example, in the case where the phone page of a called party is to be accessed, the calling party may input the URI of the called party and, indeed, this input may substitute for the normal input of a telephone number. A leading input string (for example “999”) may be used to identify the input as an URI. As regards the input means, where a user only has a standard 12 key telephone, input of host domain names and other URI elements requiring alpha characters, will need to be done using one of the standard techniques for alpha input from a phonepad (such techniques are already used, for example, to enable a calling party to “spell” out the name of the called party). It would also be possible to provide users with a full alphanumeric keypad to facilitate URI input.

Computation

Service resource access over the Internet could be restricted to a set of dialed numbers from which it was possible to compute a corresponding URI; in this case, this computation would be the responsibility of block **47**.

Association Table Lookup

Probably the simplest implementation for the block **47** is as an association table (either in memory or held on database disc store **48**) associating a URI with the UI part of each resource code. A potential problem with this approach is that a service resource may be required for a called party number on the other side of the world which implies a rigorous update regime between PSTN operators worldwide in order to keep the association table up-to-date. (Note that the same implication is not necessarily applicable in respect of marking the called-party number as one required to trigger a service request, since the number may be arranged to be one of a group of numbers all triggering an appropriate service request, in a manner similar to 800 numbers).

DNS-Type Lookup

An alternative lookup solution is to use a hierarchically-structure-d distributed database system, similar to (or even part of) the Domain Name System (DNS) of the Internet, in order to resolve the UI part of a resource code to a corresponding URI. This approach, which will be described in more detail below, would typically involve databases maintained by each PSTN operator for its numbers with which URIs are associated. These databases would be accessible by all

16

PSTNs through a network such as the Internet with resolution requests being pointed to the appropriate database in a manner similar to the Domain Name System. In this case, the block **47** is constituted by an appropriate resolution program arranged to request UI resolution over the Internet through interface **44**.

Before describing a DNS-type lookup implementation for the URI determination block **47**, some further general comments are appropriate. Whatever method is used to determine the URI, certain simplifications are possible if limited constraints are placed on the URIs permitted. In particular, it is not necessary to determine all components of an URI in the following cases:

(i) A part of the URI path component can be made standard for all service resources, this standard part being simply added by the block **47** once the rest of the URI has been determined. For example, where a roaming number is to be looked up, it may by convention always be held in a file “roam” in a subdirectory “tel” of a subscriber’s directory on a particular server. In this case the URI host component and the subscriber-unique part of the path component are first determined and then the remaining path part “/tel/roam” is added.

(ii) The URI path component can be arranged to be the same as a predetermined part of the resource code, the block **47** needing only to determine the host component and then add the path. For example, it may be agreed that the path must always end with the telephone number concerned, or sufficient of the terminating digits to have a high probability of uniqueness on the host machine. The path may also include standard components to be added by block **47**.

(iii) Blocks of telephone numbers may have their corresponding service resources located on the same host server so that it is only necessary to use a part of the telephone number to determine the host component of the URI; in this case, the path component can conveniently include all or part of each telephone number. This situation implies a tight degree of control by the telephone operators and does not offer the telephone user the freedom to choose the host server on which user places their phone page.

Another general point worthy of note is that however the URI is determined, the host component of the URI may be provided either in the form of a host domain name or a host IP address. Where the host is identified by a domain name, then a further resolution of URI host name to IP address will subsequently be carried out in standard manner by interface **44** using the Domain Name System of the Internet. This further resolution can be avoided if the host identity is directly provided as an IP address.

Where a database lookup is used to provide the number to URI translation, this database may be independent of, or combined with, a customer database containing other customer-related information. Factors affecting this choice include, on the one hand, the possible desirability of having the number-to-UI translation information widely available, and on the other hand, the possible desirability of restricting access to other customer-related information.

DNS-Type URI Lookup

A DNS-type lookup implementation for the URI determination block **47** will now be described in some detail for the case where the UI part of the resource code is a telephone number and there are no constraints on the URI, thereby requiring both the full host and path components of the URI to be returned by the lookup. A key part of the overall process is the formation of the equivalent of a host domain name from the telephone number of interest; this domain-name equivalent is then resolved into a corresponding URI by a lookup

US 8,204,046 B2

17

mechanism which in the present example is identical to that employed by the DNS (indeed, the lookup mechanism may be incorporated into the DNS though it can also be independently implemented).

The nature of the DNS has already been described above with reference to FIG. 3 when the term “DNS-type” system was also introduced. For convenience in the following a DNS-type system organised to provide a telephone number to URI translation facility will be referred to as a “Duris” system (standing for “DNS-type URI Server” system).

The basic principles surrounding operation of a Duris system are:

every telephone number can be turned into a host domain name (the name space containing such host domain names for the telephone numbers of interest is referred to below as the “telname space”); and

for every host domain name in the host domain space there is a Registration Record held by the Duris system containing the corresponding URI.

Thus, an input telephone number forming, in the present case, the UI part of a resource code 54 (see FIG. 11), is first parsed to form a domain name (step 120) and then passed to the Duris system (illustrated in FIG. 11 as formed by the DNS itself) to retrieve the RR with the corresponding URI (step 121). Following on from the URI lookup, if the URI returned has its host component as a domain name, the DNS is next used to derive the host IP address (step 122); this step is, of course not needed if the host component is stored as an IP address in the RR. The URI is then used to make a resource request to the appropriate server, passing any RRI part of the resource code 54 (step 123).

There are a number of possibilities at the top level as to how a Duris system could be implemented:

(a) Independent of the DNS. In this option, the telname space constitutes the entire name space to be managed with the root of the telname space being the “.” name space root (see FIG. 12A where the telname space is shown hatched). In this case, the Duris system is independent of the DNS itself. The Duris system could, of course, use the same basic infrastructure as the DNS (that is, the Internet) or an entirely separate network. Where the telname space comprises all the domain names corresponding to all public telephone numbers worldwide, parsing a full international telephone number would give a fully qualified domain name. Of course, the telname space could be a much smaller set of names such as those derived from internal extension numbers within a company having worldwide operations.

(b) Unfragmented Telname Space within the DNS. In this option, the telname space is a domain of the DNS name space and the Duris system is provided by the DNS itself. Thus, where the telname space comprises all domain names derived from public telephone numbers worldwide, the telname space could be placed within the domain of the ITU, in a special subdomain “tel”, the root of the telname space then being “tel.itu.int.” (see FIG. 12B where again, the hatched area represents the telname space). The responsibility for administering the domain “tel.itu.int.” would then lie with the ITU. With this latter example, to form a fully qualified domain name from an input telephone number, after the number has been parsed to form the part of the domain name corresponding to the structuring within the telname space, the tail “tel.itu.int.” is added. The fully qualified domain name is then applied to the DNS and the corresponding RR record, holding the required URI, is retrieved. As a further example, the telname space could be all name derived from internal extension numbers within Hewlett-Packard in which case the root

18

of the telname space would be “tel.hp.com.” and Hewlett-Packard would be entirely responsible for managing this domain.

(c) Fragmented Telname Space within the DNS. In this option, the telname space is split between multiple domains of the DNS name space and the Duris system is provided by the DNS itself. Thus where the telname space comprises all domain names derived from public telephone numbers worldwide, the telname space could be split between respective “tel” subdomains of each country domain; thus, as illustrated in FIG. 12C, the part of the telname space corresponding to French telephone numbers would have a root of “tel.fr.” and the part of the telname space corresponding to UK telephone numbers would have a root of “tel.uk.”. The responsibility for administering each “tel” subdomain would then lie with each country. With this latter example, to form a fully qualified domain name from an input telephone number, the part of the telephone number following the country code is parsed to form the part of the domain name within a country ‘tel’ subdomain and then a host domain name tail is added appropriate for the country concerned. Thus for a French telephone number, the “33” country code is stripped from the number prior to parsing and used to add a tail of “tel.fr.”. The tail appropriate to each country can be stored in a local lookup table. As a further example, two commercial organisations (X company and Y company) with respective DNS domains of “xco.com.” and “yco.com.” may agree to operate a common Duris system with a telname space split between “tel.xco.com.” and “tel.yco.com.”. In this case, any Y company telephone number input from X company will be parsed to a fully qualified domain name terminating “tel.yco.com.” and vice versa.

Consideration will next be given to the parsing of a telephone number into a domain name—in other words, where to insert the “.” characters into the number to provide the structuring of a domain name. Of course, as already explained, telephone numbers are hierarchically structured according to each country’s numbering plan. Thus one approach would be to follow this numbering plan structuring in dividing up a telephone number to form a domain name. By way of example, the telephone number “441447456987” which is a UK number (country code “44”) with a four digit area code (“1447”) and six digit local number (“456987”) could be divided to form a domain name of 456987.1447.44 (note that the reversal of label order occasioned by the fact that the DNS labels are arranged least significant first). If the telname space is a subdomain of the DNS with a placement as illustrated in FIG. 12B, the fully qualified domain name derived from the telephone number would be:

456987.1447.44.tel.itu.int.

There are however, difficulties inherent with trying to match the numbering plan hierarchy when parsing a telephone number into a host name. Firstly, in order to parse an international number correctly, it would be necessary for each entity tasked with this operation to know the structuring of each country’s numbering plan and where, as in the UK, area codes may be of differing length the required knowledge may need to take the form of a lookup table. Whilst this is not a complicated computational task, it is a major administrative nuisance as it means that each country will need to inform all others about its numbering plan and any updates. The second problem is that a six or seven digit local number is a very large domain; it would be preferable to create subdomains for performance and scaling reasons but there is no obvious way of doing this.

These problems can be overcome by giving up the restriction that the parsing of telephone number into a domain name

should match the structuring of national numbering plans. In fact, there is no strong reason to follow such a scheme as DNS servers know nothing about the meaning of the name space. It is therefore possible to parse telephone numbers using a deterministic algorithm taking, for example, 4 digits at a time to limit the size of each subdomain and making it possible to 'insert the dots' without knowing the numbering plan concerned. So long as the DNS domains and zones served by the DNS servers are created correctly it will all work.

For international numbers it would still seem appropriate to separate off the country codes and so a hybrid parsing scheme would be to parse the initial part of a dialed number according to known country codes and thereafter use a deterministic scheme (for example 3, 7 or 4, 6 or 3,3,4) to separate the digits. Of course, if a fragmented telname space is being used as illustrated in Figure UC then the country code is used to look up the host name tail and it is only the national part of the number which would be parsed.

Finally, as regards the details of how a DNS server can be set up to hold RR records with URIs, reference can be made, for example, to "DNS and BIND", Paul Albitz and Cricket Liu, O'Reilly & Associates, 1992 which describes how to set up a DNS server using the Unix BIND implementation. The type of the RR records is, for example, text.

It should be noted that DNS labels should not in theory start with a digit. If this convention is retained, then it is of course a trivial exercise when parsing a telephone number to insert a standard character as the first character of each label. Thus, a 4 digit label of 2826 would become "t2826" where "t" is used as the standard starting character.

It will be appreciated that as with domain names, where an input telephone number is not the full number (for example, a local call does not require any international or area code prefix), it would be parsed into a domain name in the local domain.

The foregoing discussion of Duris system implementation, has been in terms of translating a telephone number into an URI where the telephone number forms the full UI of a resource code and the Duris system returns a full URI. It will be appreciated that the described Duris implementation can be readily adapted to accommodate the various modification discussed above regarding the form of the UI and what parts of the URI need to be looked up. For example, where there are a number of different service resources associated with a subscriber each in its own file and the required source is identified by a pic part of the resource code, then the input telephone number will be used to look up, not the full URI, but the host component and that part of the path component up to the relevant subdirectory, the pic part of the UI then being appended to identify the required resource file.

For small local Duris implementations, it may be possible to have a single server; the implementation should still, however, be considered as of a DNS type provided the other relevant features are present.

Nature of Service Resources

Turning now to a consideration of the service resources 49, how these service resources can be provisioned onto the servers 51 will be described more fully below but, by way of present example, the service resource or resources associated with a particular PSTN user (individual or organisation, whether a calling or called party) can be placed on a server 51 over the Internet from a user terminal 53 in one or more WWW pages.

Consider the simple case where the service resource is a service data item such as a telephone number (for example, an alternative number to be tried if the user's telephone corresponding to the number dialed by a calling party is busy). This

diversion number could be made the sole service resource of a phone page of the user. The phone page URI could be a URL with scheme set to HTTP in which case the GET method could be used to retrieve the diversion number. Such an arrangement is suitable if the phone page is only to be used for functional retrieval of the diversion number. However, if the diversion number is to be visually presented at a user terminal 53, then it may be desirable to accompany the number with explanatory material (this will often not be necessary as the diversion number can be arranged to be returned into an existing displayed page that already provides context information). However, where the phone page does include explanatory material as well as the diversion number, an entity only wishing to make functional use of the phone page, could be arranged to retrieve the phone page and then extract the diversion number (this would, of course, require a standard way of identifying the information to be extracted from the phone page).

An alternative and preferred arrangement for providing for both viewing and functional access to a resource requiring explanatory material for viewing, is to use an object-oriented approach to resource design. In this case, the resource object would have two different access methods associated with it, one for purely functional use of the resource and the other enabling viewing of associated explanatory material. It would then be up to the accessing entity to access the resource object using the appropriate object method.

Yet another arrangement for providing for both viewing and functional use of the diversion number, would be to provide separate resources appropriately configured for each use, each resource having its own resource code (generally, both such resources would be placed on the same phone page and in this case the UI part of each resource code would be the same).

Retrieval of a phone page for use by a human user will generally not be as time critical as retrieval for operational use by a PSTN. Thus, while for human use the scheme specified in the URL of a service resource could be HTTP, it may be advantageous for operational use to define a special "phone" scheme (access protocol) which would result in the server 51 using an optimised access routine to access the required resource (diversion number, in the current example) and respond to the accessing entity in the minimum possible time.

Besides data items, other possible types of service resource include service logic for execution in place (at the server) with the result of this execution being returned to the entity accessing the resource; service logic downloadable from the server to the accessing entity for execution at that entity; and a logging resource for logging information passed to it by the accessing entity (or simply for logging the fact that it has been accessed). It will be appreciated that the logging resource is really just a particular case of service logic executable in place.

By way of example, a service resource constituted by execute-in-place service logic can be arranged to implement time-of-day routing, the result of executing the service logic being the telephone number to which a call should be routed taking account of the time of day at the called party's location. An example of a service resource constituted by downloadable service logic is service logic for controlling calling-party option interrogation using the facilities provided by an IP. As regards the logging resource, this can be used for recording the number of calls placed to a particular number.

Where each resource has its own phone page and the resource is present only in its unembellished functional form, then the HTTP scheme can be employed for access using the GET method for both the downloadable service logic and the

US 8,204,046 B2

21

execution-in-place service logic, and the POST method for the logging resource. If it is desired to provide an explanatory material with each service resource, then any of the solutions discussed above in relation to data items, can be used.

Where more than one service resource is to be associated with a number, then each such resource can be placed on a respective phone page with its own URI. However, the preferred approach is to place all such service resources on the same page and use the RRI part of the corresponding resource codes to enable access to the appropriate resource. The accessed resource is then treated according to its form (executed if execute-in-place service logic, returned if downloadable service data or logic).

Thus if both a diversion-number service-data resource and a time-of-day execution-in-place service-logic resource are placed on the same phone page, the diversion-number resource code might have an RRI of "1" whilst the time-of-day resource code might have an RRI value of "2".

Where calling/called party options are to be included in a service resource for presentation to such party, then as already indicated, this can conveniently be done by constituting the service resource as downloadable service logic with the chosen option possibly initiating request for a follow-up service resource.

It will be appreciated that a service resource will often be of a complex type, combining service data and/or downloadable service logic and/or execute in place service logic. A particularly powerful combination is the combination of the two types of service logic where the downloadable service logic is designed to interact with execute-in-place service logic; using this arrangement, the user can be presented with complex client-server type applications.

Example Usage of Service Resource

FIG. 13 illustrates the operation of a service making use of a resource on a server 51. This service is equivalent to a "personal number" service by which a user can be accessed through a single, unchanging number even when moving between telephones having different real numbers. To achieve this, the user requiring this service (user B in the current example) is allotted a unique personal number (here referred to as the "Webtel" number of B) from a set of numbers all of which have the same leading number string to enable an SSP to readily identify a dialed number as a Webtel number. User B has a service resource 49 on a dedicated phone page on HTTP server 51, this phone page being located at a URL here identified as "URL (B phone page)". B's phone page when accessed returns the current roaming number ("B-telNb") where B can be reached. In the simplest case, B's phone page is just a single number that can be modified by B (for example, from a terminal 53) as B moves to a different phone. More likely is that B's phone page is an execute-in-place service logic providing time of day routing.

In the present example, the association between B's Webtel number and the URL of B's phone page is stored in an association table accessible to SCP 43.

Upon a user A seeking to contact user B by dialing the Webtel number of B, the telephone 40 being used by A passes a call set up request to SSP 41 (note that in FIG. 13 the bearer paths through the telephony network are shown by the thicker lines 60, the other heavy lines indicating signaling flows). SSP 41 detects the dialed number as a Webtel number and sends a service request to SCP 43 together with B's Webtel number. SCP 43 on receiving this service request initiates a service logic program for controlling translation of B's Webtel number into a current roaming number for B; in fact, in the present case, this program simply requests the resource access block 46 to access the service resource identified by

22

B's Webtel number, (that is, B's phone page 49) and return the result of this access. To this end, block 46 first 'translates B's Webtel number into the URL of B's phone page and then uses this URL to access B's phone page over the Internet (for example, using the 'phone' scheme already referred to with a method corresponding to the HTTP GET method). This results in B's current roaming number B-telNb being passed back to block 46 and in due course this number is returned to the SSP 41 which then initiates completion of call set up to the telephone 40 corresponding to B-telNb.

The FIG. 13 example related to a called-party service; it will, of course, be appreciated that the principle of accessing service resources over the Internet can be applied to all types of services, including both calling-party and called-party services and hybrids. Thus, standard 800 number services can be implemented with the dialed 800 number resulting in access to a phone page resource constituted by execute-in-place service logic that returns the most appropriate number for controlling onward call routing.

It will be appreciated that although in the FIG. 13 example the service request from the SSP was triggered by a leading number string of a dialed number, a service request may be triggered by a variety of triggers including calling-party number, called-party number, or some other user input, such triggers being possibly qualified by call setup progress (for example, called-party number qualified by a busy status or by ringing for more than a certain time).

With respect to the logging service resource mentioned above, one possible application for such a resource is in telephone voting. In this case, dialing the voting number causes the SSP picking up the call to pass a service request to SCP 43 which then contacts the appropriate logging resource over the Internet to register a vote after which the call is terminated. To minimise bottlenecks, a logging resource could be provided at a different URL for each SCP, it being a simple matter to collect and collate voting from all these logging resources over the Internet. If an SCP with Internet access is provided at every SSP, then the risk of congestion is greatly reduced.

As already noted, a user's phone page may hold multiple service resources in which case the access request from the accessing SCP needs to contain an appropriate RRI identifying the required resource.

In the event that an SCP is to provide both a traditional IN service to some users and an equivalent service using an Internet-accessed service resource to other users, then a lookup table may need to be provided in the SCP to ensure that a service request is appropriately handled; such a lookup table can conveniently be combined with the customer record database.

Once a user, such as user B, has set up one or more phone pages specifying his desired service resources (particularly service logic defining personalised services), it is clearly logical for user B to want any PSTN operator he cares to use, to access and utilise such service resources. This is possible if the Webtel-to-URI databases are available to all operators. Thus multiple operators could be set to access B's phone page or pages. If an operator declines to use B's phone pages, B can obviously choose not to use that operator (at least where that operator provides a long haul carrier service subject to user selection). The possibility therefore arises that service provision will cease to command a premium from operators, but that the provision of phone-page utilisation by an operator will become a necessary basic feature of PSTN operation.

Provisioning and Updating Service Resources

Consideration will next be given as to how the service resources **49** are provisioned to the servers **51** and subsequently updated.

So far as provisioning is concerned, two basic actions are required: firstly, the service resource must be placed on a server **51** and, secondly, the URI of the service resource must be notified to the PSTN operator along with the trigger conditions (number plus any other condition such as point in call) calling for access to the resource; if multiple resources are provided at the same URI, then the RRI values needed to retrieve the appropriate resource for a particular trigger condition, must also be notified. This notification process will be referred to hereinafter as 'registering' the service resource with the PSTN operator; registration is, of course, necessary to enable the association tables used by SCP **43** to be set up and for trigger conditions to be set in SSPs **43**. For certain services, such as that described above with reference to FIG. **13**, it is not the user that supplies the triggering number (the Webtel number in the FIG. **13** example); instead, the PSTN operator allocates an appropriate number to the user as part of the registration process.

As to the process of placing a service resource on a server **51**, how this is carried out will depend on the attitude of the PSTN operator to the possible effects of such service resources on operation of the PSTN. Where the service resource simply returns a data item to an accessing entity, then an operator may not be too concerned about possible errors (accidental or deliberate) in implementing the service resource. However, the operator will probably be much more concerned about the proper operation of any service logic that may be returned by a resource; indeed, an operator may not permit such a service resource.

Assuming for the moment that an operator has no concerns about the nature or implementation of service resources, then how a resource is placed on a server **51** will largely depend on the nature of the server concerned. For example, if a user has a computer with network access to the Internet and this computer is used as server **51**, then the user can simply load a desired resource onto the server as a WWW phone page for external access. A similar situation arises if the server is an organisation server to which the server has access over an internal LAN. In both these latter cases, loading the resource as a WWW phone page does not itself require Internet access. However, if the server **51** is one run by an external Internet service provider, then a user can arrange to download the required service resource into the user's allocated Web site space on the server; this may or may not involve Internet access. One special case of this latter scenario is where the PSTN operator provides a special server for user phone pages containing service resources.

Except where a user's own computer acts as server **51**, placing a service resource on a server will generally involve clearing one or more levels of password protection.

As regards the origin of the service resource loaded by a user onto server **51**, this may be generated by the user or, particularly where the resource includes service logic, may be provided by a third party (including the PSTN operator).

If the PSTN operator wishes to have control over the service resources **49** to avoid any adverse effects on operation of the PSTN, two approaches are possible. Firstly, the operator could require that every resource (or, possibly, a particular subset) had to be subject to a verification process before use, appropriate measures then being taken to avoid subsequent alteration of the resource by the user (except, possibly, for particular data items); in this respect, the operator could require that the resource be placed on a server under the

operator's control and to which the user had no write access (except possibly for altering particular data items, as indicated above). A second, more attractive, approach to minimising adverse effects by the service resources **49**, is for the operator to provide standard service resources to which a user could add the user's own data (and possibly make limited functional selections in case where the resource included service logic); the customised resource would then be loaded onto a server **51** controlled by the operator. This process can be conveniently implemented for a particular resource using an HTML "form" which a user could download over the WWW from the operator-controlled server. After completing the form and activating a 'commit' graphical button of the form, the entered information would be 'posted' back to the server where the information would be used to produce a customised service resource thereafter placed on the server for access over the Internet. An advantage of this approach is that registration of the service resource with the operator is simultaneously effected. (It may be noted that if registration needs to be done as a separate act from having a service resource loaded on a server, then using an HTML form is a very convenient way to implement the registration process).

From the foregoing it can be seen that whilst the provisioning process does not necessarily require information to be passed over the Internet, in many cases this will be the best solution, particularly if an HTML form exchanged over the WWW can be used to produce a customised service resource. It should be noted that producing a customised service resource using an HTML form is not limited to cases where the PSTN operator controls the server.

As regards updating service resources, there is likely to be a need to update certain data items on a fairly frequent basis (for example, roaming number). Where the PSTN operator does not place any controls on the service resources **49**, then update is a relatively simple matter, only requiring write access to the server concerned (as already indicated, this will generally involve one or more levels of password protection). However, where the PSTN operator exercises control over the service resources, for example by only permitting customisations of standard service resources, such customised resources being loaded on servers controlled by the operator), then write access to the service resource may be tightly controlled. Again, an HTML form may conveniently be used as the medium for modifying a data item in such cases; to the operator, this has the benefit of limiting the modifications possible whilst to the user, a form interface should provide a simple route to resource modification.

For more complex updates, it may be necessary to go through a process similar to that required for initial provisioning.

Particularly where the service resources are held on a server **51** controlled by the PSTN operator, resource update will generally involve communication over the Internet.

Web User Interaction

Consideration will next be given to other possible uses of the service resources held in phone pages on the servers **51**. For example, if user B's phone page contains a diversion number, then provided this phone page is read-accessible over the Internet from user A's terminal **53**, user A can use a graphical Web browser running on terminal **53** to view B's phone page and discover B's diversion number. As earlier discussed, the diversion number may be passed to user A for display in an existing visual context giving meaning to the number, or may be passed to user A with accompanying explanatory text.

A more useful example is a current roaming number service for user B. Suppose B's phone page **49** on server **51** (see

FIG. 14) is operative when accessed to return a current roaming number where B can be reached. Further suppose that user B has a Web site with several Web pages written in HTML and each page contains a graphical 'phone' button which when activated uses the GET method to access B's phone page by its URL. Now if user A whilst browsing (arrow 66) B's Web site over the WWW from user A's terminal 53, decides that he would like to call user B to discuss some item of interest, user A simply activates the phone button 65 on the currently viewed page of B. This causes B's phone page to be accessed using the HTTP request "GET URL (B Phone Page)"—see arrow 67.

B's current number to be called is then determined and passed to user A's terminal 53 (see arrow 68) where it is displayed. An explanatory text concerning the number will generally also be displayed; for example the text "Please call me at the following number:" could be displayed, this text being provided either by the HTML script associated with the phone button, or from the phone page when returning the current number. In fact, it would probably be more helpful to provide user A, not only with the current number for reaching user B, but also with all numbers where B could be reached together with the times when B was most likely to be at each number. Since this extra information is likely to be subject to frequent change, the only sensible way to provide the information is from the phone page. Thus, B's phone page not only provides the current number for reaching B, but also a text that includes numbers and times subject to change; scripting B's phone page is, of course, done in a way that ensures that variable data need only be altered in one place.

In a further example, B's phone page might include downloadable service logic for execution at user A's terminal. This is useful where choices are to be presented to a user, each choice producing a follow-up action such as fetching a further phone page. For example, the first-accessed phone page may be a family phone page giving the general telephone number for a family but also giving the user the possibility of selecting further phone information on each family member, such as a time-of-day dependent number; in this case, each family member has their own follow-up phone page.

In the above scenarios, user A has been presented with a number to call over the PSTN. User A can now pick up his standard telephone and dial the number given. In fact, a complication arises if A only has Internet access via a SLIP/PPP connection over an ordinary, non-ISDN, PSTN line since, in this case, A's telephone line is already tied up with making Internet access when gateway 90 seeks to set up a call to A's telephone; with an ISDN connection, as two channels are available, this problem does not arise. One way of overcoming this problem would be to have user A's terminal 53, after obtaining the number to call from B's phone page, automatically suspend its Internet session by storing any required state information (for example, current WWW URL being accessed) and then terminate its SLIP/PPP connection to thereby free up the telephone line. A can then telephone B. At the end of this call, A can resume the suspended Internet session, using the stored state information to return to the point where A left off to call B. An alternative approach is to operate a suitable multiplexing modulation scheme on the telephone line to A allowing voice and data to be simultaneously carried. A number of such schemes already exist. The PSTN would then need to separate the combined data and voice rams coming from A at some point and pass each to its appropriate destination (the Internet data being forwarded to the ISP providing the SLIP/PPP connection for user A and the voice stream being passed to B); of course, data and voice

traffic in the reverse direction would also need combining at some point for sending over the last leg to A's terminal.

Rather than A manually dialing B using a standard telephone, another possibility is that user A's terminal is provided with functionality enabling A to make a call over the PSTN from his terminal; this functionality generally comprises a hardware interface 70 (FIG. 14) to a telephone line and phone driver software 71 for driving the interface 70 in response to input from application software such as the Web browser 73. A could call up his phone software and enter the required number or, preferably, A need only "select" on screen the number returned from B's phone page and then pass it into A's phone software. Indeed, provided user B knew the software interface to the software 71 providing dialing functionality on A's terminal, it would be possible for B's phone page to return to A's terminal program code for automatically dialing B's number upon A confirming that he wishes to proceed with call placement. As an alternative to placing a voice call, if A's terminal is equipped with a suitable modem and controlling software, A could, instead, elect to send a fax or data to B through the PSTN either to B's ordinary number or to one specified in B's phone page as the number to be used for such transmissions. Of course, placing a call from A's terminal over the PSTN may be subject to the problem already discussed of conflict for use of the telephone line where this is not an ISDN line and A gains Internet access via a SLIP/PPP connection.

However the call is placed, if B's telephone corresponding to the number tried by A is busy, a number of possibilities exist. Thus if B has a phone page that specifies a diversion number, and B has registered this service resource with the PSTN, then the diversion number should be automatically tried by the PSTN. However, if the diversion number resource has not been registered with the PSTN a busy signal will be returned to A. Where A has placed the call through a standard telephone, A must now decide how to proceed and A may elect either to give up or to refer again to B's phone page to look up the diversion number and redial using this number. If A placed the original call using his terminal 53 then the latter can be programmed to detect the return of a busy signal and then automatically look up B's diversion number and redial using this number. This functionality can be included in service logic downloaded from B's phone page and run on A's terminal.

If A had to terminate his Internet session in order to free up the telephone line for voice use, then referring back to B's phone page requires a new Internet session to be started (in fact, this inconvenience could be avoided if B's diversion number were passed to A's terminal at the time the original number to be dialed for B was supplied).

The service resource accessed on B's phone page upon B's telephone being busy may, of course, be more complex than just a diversion number. In particular, user A may be presented with a range of options including, for example, B's fax or voice mailbox number, the selection of an option potentially initiating the running of appropriate accessing software. Another possible option would be for A to leave B a call back message using a form downloaded from B's phone page upon this option being chosen; the completed form would be posted back to server 51 and logged for B to check in due course.

Of course, it may arise that user A wishes to access B's phone page to find out, for example, B's current roaming number, but user A does not know the URI of B's Web site and only has B's Webtel number. A could just call B through the PSTN in which case the translation of B's Webtel number to roaming number would be automatically effected (assuming B is still registered for this service); however, A may not wish

US 8,204,046 B2

27

to call B straight away, but just note his current roaming number. In order to solve A's problem, the Webtel-to-URI association tables previously described are preferably made accessible on the Internet at a known address (for example, at a known Web site). All that A need now do is to access this Web site passing B's Webtel number; B's phone page URI will then be returned to A who can then use it to access B's phone page. This process can, of course be made automatic from the point when A sends B's Webtel number to the association-table Web site.

Internet/PSTN Call Interface

In the FIG. 14 scenario, A's access to the PSTN was through a standard telephone interface even though the actual form of A's telephone differed from standard by being integrated into A's computer terminal 53. FIG. 15 illustrates a situation where A, after being supplied with B's current roaming number as in the FIG. 14 case, calls B via a route that starts out over the Internet and then passes through a user network interface 80 into the PSTN. Interface 80 is arranged to convert between ISDN-type telephone signaling on the PSTN and corresponding signaling indications carried across the Internet in IP packets; in addition, interface 80 transfers voice data from IP packets onto trunk 60 and vice versa.

Thus, upon A initiating a call to B, Internet phone software 81 in A's terminal sends call initiation signaling over the Internet to interface 80, the address of which is already known to A's terminal. At interface 80, the signaling is converted into ISDN-type signaling and passed to SSP 41. Call set up then proceeds in the normal way and return signaling is transferred back through interface 80, over the Internet, to the software 81 in A's terminal. This software passes call setup progress information to the WWW browser 73 for display to A. Upon the call becoming established, A can talk to B through his telephone and A's voice input is first digitised in phone hardware interface 83 and then inserted into IP packets by software 81 to traverse the Internet to interface 80 (see arrow 84); voice traffic from B follows the reverse path.

IN services can be provided to this call by SCP in response to a service request from an SSP 41. Thus, if B's phone is busy, and B is registered for call diversion, SCP 43 on receiving a service request will access B's appropriate phone page for call diversion and retrieve the diversion number. If SSP 41 is not set to initiate a service request on B's telephone being busy, the busy indication is returned to A's terminal where it can be handled in the manner already described with reference to FIG. 14.

In fact, interface 80 can be provided with functionality similar to an SSP to set trigger conditions and generate a service request to SCP 43 on these conditions being satisfied.

Third-Party Call Setup Gateway

FIG. 16 illustrates a further arrangement by which A can call B after receiving B's current roaming number. In this case, a third-party call set-up gateway 90 is provided that interfaces both with the Internet 50 and with an SSP 41. Conveniently, gateway 90 can be co-located with SCP 43 (though this is not essential). Gateway 90 has the capability of commanding SSP 41 to set up a call between specified telephones.

Thus, upon A wishing to call B, a third-party call setup request is sent from A's terminal over the Internet to the gateway 90 (see arrow 91). This setup request includes A's telephone number and B's current roaming number. Gateway 90 first attempts to setup the call to A's telephone (which should generally succeed) and thereafter to set up the call to B's identified telephone. Once the call is setup, A and B communicate in standard manner across the PSTN.

28

If B's phone had been busy, then any of the previously described scenarios may ensue.

Gateway 90 can also be arranged to make service requests to SCP 43 upon predetermined trigger conditions being satisfied. Thus, gateway 90 might be set to pick up the busy condition on B's telephone and initiate a service request to SCP 43 for a diversion number. However passing the busy indication back to A's terminal via gateway 90 is preferred because of the flexibility it gives A regarding further action.

As already generally discussed in relation to FIG. 14, a complication arises if A only has Internet access via a SLIP/PPP connection over an ordinary, non-ISDN, PSTN line since, in this case, A's telephone line is already tied up with making Internet access when gateway 90 seeks to set up a call to A's telephone. The solutions discussed in respect of FIG. 14 (termination of Internet session; multiplexing voice and Internet data on same telephone line) can also be used here. An alternative approach both for FIG. 14 and for FIG. 16 scenarios is possible if user A's terminal can handle a voice call as digitised voice passed over the Internet. In this case, the voice call can be placed through an interface 80 of the FIG. 15 form, and the voice traffic and the Internet communication with the B's phone page and/or gateway 90 are both carried in Internet packets passed over the SLIP/PPP connection to/from A's terminal 53 but as logically distinct flows passed to separate applications running on terminal 53.

It may be noted that the third party call setup request made by A's terminal to gateway 90 could equally have been made by service logic held in B's phone page and executed by server 51 (such an arrangement would, of course, require A's telephone number to be passed to B's phone-page service logic and this could be arranged to occur either automatically or through a form presented to user A at terminal A and then posted back to server 51).

It may also be noted that the interface 80 of FIG. 15 and the gateway 90 of FIG. 16 provide examples of service requests being passed to the service control subsystem by entities other than SSPs 41.

WWW-Based "FreePhone" (800 Number) Services

It is possible to implement a "FreePhone" or "800 number" type of service using a combination of the WWW and the PSTN. As will be seen from the following description of such a service with reference to FIG. 17, a WWW/PSTN implementation does not necessarily rely either on transferring call charges from the calling to called party or on the use of a special "800" number, two characteristics of standard "Free-phone" schemes. The WWW/PSTN implementations do, however, possess the more general characteristic of placing an enquiring party and the party to whom the enquiry is directed, in telephone contact at the expense of the latter party.

In the FIG. 17 arrangement, a user D such as a large department store has a website on a server 51; for the sake of simplicity, it will be assumed that the server is under the control of user D who has direct computer access to the server over line 125. D's Website may, for example, contain many catalogue-like Web pages illustrating goods offered for sale by D. In addition, D has a freephone page 124 for handling enquiries placed on a freephone basis; the URL of this page is associated with a "freephone" graphical button 122 placed on each of the Website catalogue pages.

Suppose user A at terminal 53 is browsing D's Website, looking at the catalogue pages (arrow 121). If A sees an item of interest and wishes to make an enquiry to D about this item, then A can activate at terminal 53 the graphical freephone button 122 associated with the catalogue page concerned. This activation causes code embedded in the catalogue page

currently loaded in A's terminal to prompt the user to enter their telephone number and, optionally, their name, after which an HTTP request is sent to D's freephone page using the POST method and enclosing the entered data (arrow 123). D's freephone page on receiving this request executes service logic to enter a new enquiry (including A's name and telephone number) in an enquiry queue 127 maintained in an enquiry control system 126. In the present example, the enquiry control system is connected to the server 51 via line 125, externally of the Internet; however, it would also be possible to have server 51 communicate with the enquiry control system through the Internet and, indeed, this may be the most practical arrangement where D's Website is on an ISP server rather than on a server controlled by D. In fact, the code run in A's terminal upon activation of the freephone graphical button 122 could be arranged to directly forward the enquiry request to the enquiry control system over the Internet rather than passing it back through the server 51.

The enquiry control system 126 manages enquiries passed to it to ensure that they are dealt with in an ordered manner. The system 126 on receiving a new enquiry preferably estimates approximately how long it will be before the enquiry is dealt with, this estimation being based on the number of currently queued enquiries and the average time taken to handle an enquiry. This estimation of waiting time is passed back via server 51 to user A in the response to the POST request message.

The enquiry control system 126 looks after the distribution of enquiries to a number of agents each of which is equipped with a telephone 40 and a display 129. A's enquiry will be dealt with as soon as it reaches the head of the queue 127 and there is an agent detected as available to handle the enquiry (thus, for example, the system may be arranged to detect when an agent's telephone goes on hook). When these conditions are met, a distribution and setup control unit 128 takes A's enquiry and displays A's name and telephone number on the display 129 of the available agent (for clarity, herein referenced as agent D'); if user D keeps a database on D's past customers or credit rating data, then unit 128 will also look for and display any such further information known about A. At the same time, unit 128 makes a third-party call setup request (arrow 130) over the Internet to gateway 90 asking for a call to be set up between the telephone of the available agent D' and the telephone of user A, both telephones being identified by their respective numbers. If both D' and A pick up the call, the enquiry then proceeds, the cost of the call being paid for by D as it is D that originated the call over the PSTN. If, for whatever reason, the call remains incomplete (for example, unanswered by A) for a predetermined timeout period, then unit 128 can be arranged to automatically pass on to the next enquiry at the head of the queue 127.

It would, of course, be possible to dispense with having the unit 128 request call setup through gateway 90 and either have the agent D' dial A's number manually or have unit 126 initiate auto-dialing for D' telephone (agent D' having, for example, a computer-integrated telephone similar to that of A's in FIG. 14). The advantage of these approaches is that the existing PSTN could be used without adaption and without any service installation, in implementing the WWW-based freephone service.

As discussed in relation to FIGS. 11 and 13, a complication arises in placing a call to A if A only has Internet access via a SLIP/PPP connection over an ordinary, non-ISDN, PSTN line since, in this case, A's telephone line is already tied up with making Internet access when user D tries to set up a call to A's telephone. The solutions discussed in respect of FIGS. 11 and 13 can also be used here (termination of Internet

session; multiplexing voice and Internet data on same telephone line; and placing the call over the Internet to A's terminal). With respect to the solution based on termination of the Internet session, such termination could be delayed until A's enquiry was about to be dealt with; however, to do this, it would be necessary to provide feedback from the control system 126 over the Internet to A's terminal 53 and to associate this feedback with code for bringing about Internet-session termination. One way to achieve this would be to have the response message sent by server 51 in reply to the original POST request message from A, include a correlation code; any subsequent feedback from system 126 passed to A would also include this code (server A having also passed the code to control system 126) thereby allowing A's terminal to correctly identify this feedback. In fact, the same mechanism could be used to provide user A with updates on how much longer user A is likely to be waiting to be called back, this mechanism being usable independently of whether or not there was a conflict problem for use of A's telephone line.

Where user A only has a telephone 40 and no terminal 53, it is still possible to utilise the basic structure of FIG. 17 to provide a freephone service for user A without resorting to the complexity of call charge transfer. More particularly, A would dial a special number for user D's freephone service (typically an 800 number), and the SSP 41 would recognise this special number in standard manner and make a service request to SCP 43 including both this special number and A's number. SCP 43 would then ascertain D's freephone-page URL by doing a number-to-URL translation and access D's freephone page using a POST-method HTTP request similar to request 123. Once this request had been registered as an enquiry by D's freephone page 124, the latter could send a response to SCP 43 asking it to play an announcement such as "Your freephone enquiry has been registered; please hang up and you will be contacted shortly". This announcement could be played to A by an IP in standard manner. A would then hang up and be ready to receive a call from D.

A significant advantage of the above freephone schemes using WWW, is that user D is not running up charges for use of the PSTN during periods when an enquiry is enqueued, waiting to be handled.

Variants

Many variants are, of course, possible to the above-described arrangements and a number of these variants are described below.

Distributed Processing Environment. As is illustrated in FIG. 18, the SCP 43 may access the HTTP servers 51 through a distributed processing environment, DPE 98, at least logically separate from the Internet. Preferably in this case the servers 51 are controlled by PSTN operators and are thus restricted in number.

Service Resources on DNS-Type Servers. In the foregoing examples, the service resource items have been placed on servers 51 connected to the Internet and a desired service resource has then been accessed over the Internet by the service control subsystem of the PSTN, and/or by Internet users, through the use of an URI derived from a resource code that identifies the desired service resource item. In a preferred arrangement for deriving the URI from a resource code in the form of a telephone number, all or part of the telephone number concerned was parsed into domain name form and then resolved into an URI using a DNS-type distributed database system that, indeed, could be integrated into the DNS itself (see FIGS. 11 and 12, and related description). In fact, it would be possible to place service resource items directly in Registration Records held by a DNS-type distributed database system so that instead of the parsed telephone number

US 8,204,046 B2

31

being resolved to an URI which is then used to access the required resource, the parsed telephone number is directly resolved to the required service resource item. The mechanism employed in this process is exactly as already described for resolving a parsed telephone number into an URI. The DNS-type distributed database system used for this would preferably be one accessible over the Internet or the DNS itself so as to provide access to the service resource items for Internet users as well as for the service control subsystem of the PSTN (in the same manner as described above with reference to FIG. 18, the DNS-type servers holding the service resource items may be accessible to the service control subsystem by a network other than the Internet). Whilst the placing of service resource items in RRs held on DNS-type servers may not be suitable for all types of service resource items, it is suitable for items such as telephone numbers that do not change frequently. Thus, a suitable usage is to provide number portability; in this case, a dialed personal number triggers a lookup in the DNS-type system with all or part of the personal number being first parsed and then applied to the DNS type system to return a current number for call routing. All dialed numbers could be treated as personal numbers or simply a subset of such numbers, this subset comprising numbers that are readily identifiable as personal numbers by, for example, local lookup at an SSP or the presence of a predetermined leading digit string. The general concept of parsing a telephone number (or similar number) in whole or in part to form a domain name for resolution in a DNS-type distributed database system can be used for the retrieval of other items of information besides URIs and service resource items.

Feedback Mechanisms. In discussing the WWW-based freephone arrangement of FIG. 17, it was mentioned that user A could be supplied with feedback on the likely length of waiting time before A would be called back. This is one example of using the Internet to provide a feedback path for a potential or actual telephone user. Another example was provided in relation to FIG. 16 where the progress of call setup was reported back by the call setup gateway to user A's terminal. In fact, generally where a user is known to be using a terminal actively on the Internet the opportunity arises to provide the user with feedback on the progress of call setup through the telephone system. In order to do this, it is of course necessary to ensure that the feedback can be passed to the appropriate application running on terminal A and this will generally require the application to have made appropriate linking information available. As well as call setup progress information, other information can also be feedback for example during a call holding period. Thus, for example, a special server can be provided on the Internet holding multimedia clips or even videos that could be output to user A during a call holding period.

In the described arrangements, the servers 51 have held service resource items concerned primarily with call setup control. It may be noted that in a somewhat different application, Internet servers could be arranged to hold data that could be accessed from the telephone system in response to a user-initiated telephone request and returned to that telephone user. Such a service would be provided, for example, in response to an SSP triggering a service request upon a particular telephone number being input, the service request prompting an SCP to cause an intelligent peripheral to access a particular Internet server (not necessarily an HTTP server) and retrieve the required data for return to the calling party. The intelligent peripheral may include a text-to-voice converter for replaying the data vocally to the user.

32

One further feedback process is also worthy of note, in this case in relation to service resource items themselves. By way of example, a telephone user G may subscribe to a service by which calls passed through to G's telephone are to be separated by a minimum of X minutes, X being user settable. To implement this service, G has a phone page on a server 51 that includes a "busy" status indication. Upon termination of a successful call to G, G's local SSP triggers the sending of a message by the associated SCP over the Internet to G's phone page. This message causes G's busy indication to be set to indicate that G is busy; the message also starts a timer which times out after a period X and causes the busy status indication to be reset. A call attempt to G will either be rejected at G's SSP because G's line is genuinely busy or will trigger the SSP to enquire via the SCP whether G's phone-page busy status indication is set. If the busy status indication is set (which it will be during the period X following termination of a successful call) the call attempt is rejected whereas if the busy status indication is in its reset condition, the call attempt is allowed to proceed.

By placing the busy status indication mechanism on G's phone page, it is possible to arrange for G to be able to easily change the value of X.

More General Variants. Whilst the service control subsystem of the PSTN has been embodied as an SCP in the foregoing examples, it will be appreciated that the functionality of the service control subsystem could be provided as part of an SSP or in an associated adjunct. Furthermore, the triggering of service requests can be effected by equipment other than SSPs, for example by intercept boxes inserted in the SS7 signaling links.

It will be appreciated that the term "Internet" is to be understood to include not only the current specification of the TCP/IP protocols used for the Internet and the current addressing scheme, but also evolutions of these features such as may be needed to deal with isochronous media. Furthermore, references to the WWW and the HTTP protocol should equally be understood to encompass their evolved descendants.

The present invention can also be applied to telephone systems other than just PSTNs, for example to PLMNs and other mobile networks, and to private systems using PABXs. In this latter case, a LAN or campus-wide computer network serving generally the same internal users as the PABX, will take the role of the Internet in the described embodiments.

Furthermore, the present invention has application where any switched telecommunication system (for example, a broadband ATM system) requires service control and a computer network can be used for the delivery of service resources to the service control subsystem of the telecommunication system.

The invention claimed is:

1. An apparatus for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, said apparatus comprising:

a processor; and

a non-transitory computer-readable storage medium storing computer-executable instructions, which when executed cause the apparatus to:

US 8,204,046 B2

33

apply said process to said number string identifying the target entity to form the related domain name, and supply the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity.

2. The apparatus according to claim 1, wherein the apparatus is further configured to use the retrieved URI to access said target entity.

3. The apparatus according to claim 2, wherein the target entity is an information entity, the apparatus being further configured to access the target entity to retrieve information.

4. The apparatus according to claim 2, wherein the target entity is a logging entity, the apparatus being further configured to access the target entity to log information.

5. The apparatus according to claim 1, wherein the URI of the target entity comprises an access scheme and host address.

6. The apparatus of claim 1, wherein the URI is a universal resource name (URN).

7. The apparatus of claim 1, wherein the URI indicates a location of a phone page.

8. The apparatus of claim 7, wherein the phone page comprises a hypertext markup language (HTML) page.

9. The apparatus of claim 7, wherein the phone page comprises a page of the World Wide Web of the Internet.

10. The apparatus of claim 1, wherein the URI is an Internet-resolvable indicator of a location of a phone page.

11. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name; and

supplying the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity.

12. The non-transitory computer readable medium according to claim 11, wherein the number string identifying the target entity is of telephone-number form.

13. The non-transitory computer readable medium according to claim 11, wherein the URI of the target entity comprises an access scheme and host address.

14. The non-transitory computer readable medium according to claim 11, wherein the target entity is an information entity, the method further comprising using the URI to access the target entity to retrieve information.

15. The non-transitory computer readable medium according to claim 11, wherein the target entity is a logging entity, the method further comprising accessing the target entity to log information.

16. The non-transitory computer-readable medium of claim 11, wherein the URI is a universal resource name (URN).

17. The non-transitory computer-readable medium of claim 11, wherein the URI indicates a location of a phone page.

18. The non-transitory computer-readable medium of claim 17, wherein the phone page comprises a hypertext markup language (HTML) page.

19. The non-transitory computer-readable medium of claim 17, wherein the phone page comprises a page of the World Wide Web of the Internet.

20. The non-transitory computer-readable medium of claim 11, wherein the URI is an Internet-resolvable indicator of a location of a phone page.

34

21. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name; and

supplying the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

22. The non-transitory computer readable medium according to claim 21, wherein the said number string identifying the target entity is of telephone-number form.

23. The non-transitory computer readable medium according to claim 21, wherein the URI of the target entity comprises an access scheme and host address.

24. The non-transitory computer readable medium according to claim 21, wherein the target entity is an information entity, the method further comprising using the URI to access the target entity to retrieve information.

25. The non-transitory computer readable medium according to claim 21, wherein the target entity is a logging entity, the method further comprising using the URI to access the target entity to log information.

26. The non-transitory computer-readable medium of claim 21, wherein the URI is a universal resource name (URN).

27. The non-transitory computer-readable medium of claim 21, wherein the URI indicates a location of a phone page.

28. The non-transitory computer-readable medium of claim 27, wherein the phone page comprises a hypertext markup language (HTML) page.

29. The non-transitory computer-readable medium of claim 27, wherein the phone page comprises a page of the World Wide Web of the Internet.

30. The non-transitory computer-readable medium of claim 21, wherein the URI is an Internet-resolvable indicator of a location of a phone page.

31. An apparatus, comprising:

a network interface;

a processor; and

a database storing a plurality of resource records that each provides a mapping from a domain name to a uniform resource identifier (URI) for locating a target entity associated with the domain name, at least a portion of each domain name being in the form of a number string that has been parsed into plural domain-name labels, wherein the processor is configured to, in response to an input identifying one of the domain names, output an associated one of the URIs stored in the database to the network interface.

32. The apparatus according to claim 31, wherein the database comprises the domain name system of the Internet.

33. An apparatus for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, said apparatus comprising:

US 8,204,046 B2

35

a processor; and

a non-transitory computer-readable storage medium storing computer-executable instructions, which when executed cause the apparatus to:

apply said process to said number string identifying the target entity to form the related domain name, by parsing the number string with a deterministic algorithm, and

supply the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity.

34. An apparatus for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, said apparatus comprising:

a processor; and

a non-transitory computer-readable storage medium storing computer-executable instructions, which when executed cause the apparatus to:

apply said process to said number string identifying the target entity to form the related domain name, wherein the number string comprises a plurality of digits, and wherein said applying the process comprises parsing the number string such that the domain name also comprises the plurality of digits, and

supply the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity.

35. An apparatus for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, said apparatus comprising:

a processor; and

a non-transitory computer-readable storage medium storing computer-executable instructions, which when executed cause the apparatus to:

apply said process to said number string identifying the target entity to form the related domain name, wherein the number string comprises a plurality of digits, and wherein said applying the process comprises reversing at least some of the digits to form the domain name, and

supply the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity.

36. An apparatus for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, said apparatus comprising:

36

a processor; and

a non-transitory computer-readable storage medium storing computer-executable instructions, which when executed cause the apparatus to:

apply said process to said number string identifying the target entity to form the related domain name, by parsing each of a plurality of portions of the number string into a different one of a plurality of portions of the domain name, and

supply the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity.

37. An apparatus for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, said apparatus comprising:

a processor; and

a non-transitory computer-readable storage medium storing computer-executable instructions, which when executed cause the apparatus to:

apply said process to said number string identifying the target entity to form the related domain name, and

supply the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity,

wherein the apparatus is further configured to send data to the target entity identified using the URI.

38. An apparatus for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, said apparatus comprising:

a processor; and

a non-transitory computer-readable storage medium storing computer-executable instructions, which when executed cause the apparatus to:

apply said process to said number string identifying the target entity to form the related domain name, and

supply the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity,

wherein the apparatus is further configured to receive the number string and data from a phone, and to send the data to the target entity identified using the URI.

39. An apparatus for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, said apparatus comprising:

US 8,204,046 B2

37

a processor; and

a non-transitory computer-readable storage medium storing computer-executable instructions, which when executed cause the apparatus to:

apply said process to said number string identifying the target entity to form the related domain name, and supply the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity,

wherein the apparatus is further configured to send a message to the target entity identified using the URI.

40. An apparatus for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, said apparatus comprising:

a processor; and

a non-transitory computer-readable storage medium storing computer-executable instructions, which when executed cause the apparatus to:

apply said process to said number string identifying the target entity to form the related domain name, and supply the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity,

wherein the apparatus is further configured to set up a voice call to the target entity identified using the URI.

41. An apparatus for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, said apparatus comprising:

a processor; and

a non-transitory computer-readable storage medium storing computer-executable instructions, which when executed cause the apparatus to:

apply said process to said number string identifying the target entity to form the related domain name, and supply the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity,

wherein the apparatus is further configured to contact the target entity identified using the URI.

42. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing, using a deterministic algorithm, at least a portion of the number string into at least a part of said domain name; and

supplying the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity.

43. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

38

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein the number string comprises a plurality of digits, and wherein said parsing comprises parsing the number string such that the domain name also comprises the plurality of digits; and

supplying the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity.

44. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein the number string comprises a plurality of digits, and wherein said parsing comprises reversing at least some of the digits to form the domain name; and

supplying the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity.

45. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein said parsing comprises parsing each of a plurality of portions of the number string into a different one of a plurality of portions of the domain name; and

supplying the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity.

46. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity; and sending data to the target entity identified using the URI.

47. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

receiving a number string and data from a phone;

forming, from the number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity; and sending the data to the target entity identified using the URI.

US 8,204,046 B2

39

48. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity; and sending a message to the target entity identified using the URI.

49. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity; and setting up a voice call to the target entity identified using the URI.

50. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity; and contacting the target entity identified using the URI.

51. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein said parsing comprises parsing the number string with a deterministic algorithm; and supplying the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

52. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein the number string comprises a plurality of digits, and wherein said parsing comprises parsing the number string such that the domain name also comprises the plurality of digits; and

supplying the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

53. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a

40

portion of the number string into at least a part of said domain name, wherein the number string comprises a plurality of digits, and wherein said parsing comprises reversing at least some of the digits to form the domain name; and

supplying the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

54. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein said parsing comprises each of a plurality of portions of the number string into a different one of a plurality of portions of the domain name; and supplying the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

55. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity;

sending data to the target entity identified using the URI.

56. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

receiving a number string and data from a phone;

forming, from the number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity;

sending the data to the target entity identified using the URI.

57. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity;

sending a message to the target entity identified using the URI.

58. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity; and

US 8,204,046 B2

41

setting up a voice call to the target entity identified using the URI.

59. A non-transitory computer readable medium storing computer-executable instructions for performing a method, the method comprising:

forming, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity; and

contacting the target entity identified using the URI.

60. A method for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, the method comprising:

applying, by at least one computing device, said process to said number string identifying the target entity to form the related domain name; and

supplying, by the at least one computing device, the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity.

61. The method of claim 60, wherein the URI is a universal resource name (URN).

62. The method of claim 60, wherein the URI indicates a location of a phone page.

63. The method of claim 62, wherein the phone page comprises a hypertext markup language (HTML) page.

64. The method of claim 62, wherein the phone page comprises a page of the World Wide Web of the Internet.

65. The method of claim 60, wherein the URI is an Internet-resolvable indicator of a location of a phone page.

66. A method for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, the method comprising:

applying, by at least one computing device, said process to said number string identifying the target entity to form the related domain name, by parsing the number string with a deterministic algorithm; and

supplying, by the at least one computing device, the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity.

67. A method for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, method comprising:

42

applying, by at least one computing device, said process to said number string identifying the target entity to form the related domain name, wherein the number string comprises a plurality of digits, and wherein said applying the process comprises parsing the number string such that the domain name also comprises the plurality of digits; and

supplying, by the at least one computing device, the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity.

68. A method for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, method comprising:

applying, by at least one computing device, said process to said number string identifying the target entity to form the related domain name, wherein the number string comprises a plurality of digits, and wherein said applying the process comprises reversing at least some of the digits to form the domain name; and

supplying, by the at least one computing device, the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity.

69. A method for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, the method comprising:

applying, by at least one computing device, said process to said number string identifying the target entity to form the related domain name, by parsing each of a plurality of portions of the number string into a different one of a plurality of portions of the domain name; and

supplying, by the at least one computing device, the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity.

70. A method for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, the method comprising:

applying, by at least one computing device, said process to said number string identifying the target entity to form the related domain name;

supplying, by the at least one computing device, the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity; and

sending data to the target entity identified using the URI.

71. A method for facilitating accessing a target entity identified at least in part by a number string, in which a domain

US 8,204,046 B2

43

name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, the method comprising:

receiving the number string and data from a phone;

applying, by at least one computing device, said process to said number string identifying the target entity to form the related domain name;

supplying, by the at least one computing device, the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity; and

sending the data to the target entity identified using the URI.

72. A method for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, the method comprising:

applying, by at least one computing device, said process to said number string identifying the target entity to form the related domain name;

supplying, by the at least one computing device, the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity; and

sending a message to the target entity identified using the URI.

73. A method for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, the method comprising:

applying, by at least one computing device, said process to said number string identifying the target entity to form the related domain name;

supplying, by the at least one computing device, the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity; and

setting up a voice call to the target entity identified using the URI.

74. A method for facilitating accessing a target entity identified at least in part by a number string, in which a domain name system (DNS) of a packet-switched network stores records each associated with a corresponding domain name and holding a URI of an entity associated with the domain name, each said domain name being related to a respective number string from which it can be derived by a process including parsing at least a portion of the number string into at least a part of said domain name, the method comprising:

apply said process to said number string identifying the target entity to form the related domain name;

supply the domain name to the DNS to retrieve the URI held in the corresponding record, the URI being that of the target entity; and

contacting the target entity identified using the URI.

44

75. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name; and

supplying, by the at least one computing device, the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity.

76. The method of claim 75, wherein the URI is a universal resource name (URN).

77. The method of claim 75, wherein the URI indicates a location of a phone page.

78. The method of claim 77, wherein the phone page comprises a hypertext markup language (HTML) page.

79. The method of claim 77, wherein the phone page comprises a page of the World Wide Web of the Internet.

80. The method of claim 75, wherein the URI is an Internet-resolvable indicator of a location of a phone page.

81. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing, using a deterministic algorithm, at least a portion of the number string into at least a part of said domain name; and

supplying, by the at least one computing device, the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity.

82. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein the number string comprises a plurality of digits, and wherein said parsing comprises parsing the number string such that the domain name also comprises the plurality of digits; and

supplying, by the at least one computing device, the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity.

83. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein the number string comprises a plurality of digits, and wherein said parsing comprises reversing at least some of the digits to form the domain name; and

supplying, by the at least one computing device, the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity.

84. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein said parsing comprises parsing each of a plurality of portions of the number string into a different one of a plurality of portions of the domain name; and

US 8,204,046 B2

45

supplying, by the at least one computing device, the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity.

85. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity; and

sending data to the target entity identified using the URI.

86. A method, comprising:

receiving a number string and data from a phone;

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity; and

sending the data to the target entity identified using the URI.

87. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity; and

sending a message to the target entity identified using the URI.

88. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity; and

setting up a voice call to the target entity identified using the URI.

89. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a domain name system of a packet-switched network and receiving back from the domain name system a resource record including a uniform resource identifier (URI) of the target entity; and

contacting the target entity identified using the URI.

90. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a

46

process including parsing at least a portion of the number string into at least a part of said domain name; and

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

91. The method of claim 90, wherein said forming comprises determining a label.

92. The method of claim 90, wherein said forming comprises determining an alphanumeric label.

93. The method of claim 90, wherein said forming comprises determining a numeric label.

94. The method of claim 90, wherein the number string used in said forming is associated with a dialed number.

95. The method of claim 90, further comprising using the number string to communicate across a switched telecommunication system.

96. The method of claim 90, further comprising using data associated with the number string to route calls in a switched telecommunication system.

97. The method of claim 90, further comprising using data associated with the number string to route messages in a switched telecommunication system.

98. The method of claim 90, further comprising using data associated with the number string to route data packets in a switched telecommunication system.

99. The method of claim 90, wherein the URI is a universal resource name (URN).

100. The method of claim 90, wherein the URI indicates a location of a phone page.

101. The method of claim 100, wherein the phone page comprises a hypertext markup language (HTML) page.

102. The method of claim 100, wherein the phone page comprises a page of the World Wide Web of the Internet.

103. The method of claim 90, wherein the URI is an Internet-resolvable indicator of a location of a phone page.

104. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity; and

using data associated with the number string to route data packets in a switched telecommunication system, wherein the switched telecommunication system is logically distinct from a portion of a communication system originating data associated with the number string.

105. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity;

using data associated with the number string to route data packets in a switched telecommunication system;

provisioning at least one server with a plurality of addresses associated with the switched telecommunication system;

providing a mapping associated with the switched telecommunication system; and

US 8,204,046 B2

47

establishing a channel through the switched telecommunications system utilizing the at least one server and the mapping.

106. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity;

provisioning at least one server with a plurality of addresses associated with a switched telecommunication system;

providing a mapping associated with the switched telecommunication system in cooperation with the number string

establishing a bearer channel through the switched telecommunications system utilizing the at least one server and the mapping.

107. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein said parsing comprises parsing the number string with a deterministic algorithm; and

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

108. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein the number string comprises a plurality of digits, and wherein said parsing comprises parsing the number string such that the domain name also comprises the plurality of digits; and

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

109. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name, wherein the number string comprises a plurality of digits, and wherein said parsing comprises reversing at least some of the digits to form the domain name; and

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

110. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name,

48

wherein said parsing comprises each of a plurality of portions of the number string into a different one of a plurality of portions of the domain name; and

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity.

111. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity; and

sending data to the target entity identified using the URI.

112. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

receiving a number string and data from a phone;

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity; and

sending the data to the target entity identified using the URI.

113. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity; and

sending a message to the target entity identified using the URI.

114. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity; and

setting up a voice call to the target entity identified using the URI.

115. A method, comprising:

forming, by at least one computing device, from a number string identifying a target entity, a domain name by a process including parsing at least a portion of the number string into at least a part of said domain name;

supplying, by the at least one computing device, the domain name to a database and receiving back from that database a resource record including a uniform resource identifier (URI) of the target entity; and

contacting the target entity identified using the URI.

* * * * *

CERTIFICATE OF SERVICE

I hereby certify that on April 4, 2016, I electronically filed the foregoing Brief with the Clerk of this Court using the CM/ECF System, which will send notice of such filing to all registered CM/ECF users.

/s/ Brian D. Schmalzbach
Brian D. Schmalzbach

CERTIFICATE OF COMPLIANCE

This brief complies with the type-volume limitation of Fed. R. App. P. 32(a)(7)(B)

because:

- This brief contains 13,562 words, excluding the parts of the brief exempted by Fed. R. App. P. 32(a)(7)(B)(iii) and Fed. Cir. R. 32(b).

This brief complies with the typeface and type style requirements of Fed. R.

App. P. 32(a)(5) because:

- This brief has been prepared in a proportionally spaced 14-point Times New Roman font using Microsoft Word.

/s/ Brian D. Schmalzbach
Brian D. Schmalzbach